ROUTING AND RECORD SHEET

SUBJECT: (Optional)

Presentation of Computer Systems Performance Data

FROM:	EXTENSION	NO.
Keith E. Silliman IBM/SEB/ED/ODP	[redacted]	DATE 20 September 1983

STAT

TO: (Officer designation, room number, and building)	DATE RECEIVED	DATE FORWARDED	OFFICER'S INITIALS	COMMENTS (Number each comment to show from whom to whom. Draw a line across column after each comment.)
1. [redacted] EXO/ODP 1D0105 Hqs		21 Sept	GD	
2.				
3. C/ADP&E~~B~~/PD/OL [redacted]		9/23/83	J	3. For approval.
4.				
5. EXO/ODP 1D0105 Hqs	9/29	9/29	MA	
6.				
7.				
8.				
9.				
10.				
11.				
12.				
13.				
14.				
15.				

STAT

STAT

Fred:

Mr. Silliman is an IBM employee who has presented the same paper before to IBM groups with Agency approval. The paper covers the results of a VM performance study he did for us. He will not reveal the source of his data during the presentation. Even if he did, no harm would be done because the fact that the Agency uses VM is well known in the IBM community thanks to our active involvement in SHARE. I recommend your approval.

[redacted]
EXO/ODP

STAT

STAT

APPROVA[redacted] 9/23/83
C/ADP&EB/PD/~~OL~~ Date

1) did George already tell Keith ok?

2) Fil in IBM(?)

Yes GD

Patty – pls send copy of route sheet + cover letter to Keith Silliman, they fil. (I assume he attended his conference)

done 9/30/83 pmd

FORM 610 USE PREVIOUS EDITIONS
1-79

30 August 1983

MEMORANDUM FOR: Contracting Officer

THROUGH : Chief, Engineering Division, ODP
Executive Officer, ODP

FROM: Keith E. Silliman
IBM

SUBJECT: Presentation of Computer System Performance Data

1. I request permission to present computer system performance data obtained from Agency IBM computer systems to a group of computer performance system analysts.

2. When approved, I intend to present that data at an IBM internal meeting of the 2nd Annual VM Performance Seminar to be held in Toronto, Ontario, Canada, September 27 - 29, 1983. The paper and presentation foils will be published in an IBM Internal Only proceedings document.

3. None of the material to be presented is classified or controversial. I will discuss the technical aspects of VM computer system performance. The title of the paper is "VM Performance Analysis and Capacity Planning With Emphasis on DASD". A copy of the paper is attached along with a copy of the foils to be used for the presentation.

STAT

4. I [redacted] will be identified as an IBM employee. If any reference is made as to the source of the data, it will be identified only as from a customer account.

Keith E. Silliman

STAT

cc: [redacted]

VM PERFORMANCE ANALYSIS AND CAPACITY PLANNING WITH EMPHASIS ON DASD

by

Keith E. Silliman

IBM Corporation
Federal Systems Division
18100 Frederick Pike
Gaithersburg, Maryland 20879
Tie 372-7132
301 840-7132

This paper presents some results and conclusions from performance analysis and capacity planning work done with a large scale (3081 D), high performance (sub-second trivial response time) VM/CMS intensive computer system. We will discuss the performance analysis and capacity planning methodology, summarize some results, make observations on subsystem capacity values, and indicate future directions.

Methodology

Our approach to performance analysis and capacity planning was to characterize the capacity of the system, by subsystem component (CPU, Main Storage, and DASD), in terms of number of 'users'. (The characterization was accomplished by plotting the subsystem utilization as a function of the number of active users.) This methodology provided a sufficiently precise measure of the system environment, at a relatively low cost in system and manpower resources, for both gross level performance analysis and capacity planning functions. It indicates the percent of subsystem utilization per user, which is instrumental in system tuning when the objective is to balance subsystem capacities. Assuming that the workload mix does not significantly change, subsystem utilization can be linearly projected for capacity planning purposes, based on anticipated future user loads.

This methodology was a synthesis of the system decomposition approach of Major (IBM SYST J 20 #1 1981) and Wicks (GG22-9299, 1982) in an MVS environment and the user orientation of Tetzlaff (IBM SYST J 18 #1 1979) in a VM environment.

Results

In this environment, we see 2.0 to 2.5 logged users per active user. (Definitions of the measures discussed can be obtained from the VMAP documentation.) This relatively low ratio of logged to active users occurs because users are not allowed to leave a logged on terminal unattended. The logged to active ratio is generally indicative of user response, with the lowest ratio (approximately 2 logged to 1 active) giving the best response. A fully utilized (190 to 200 percent) 3081 D processor supports between 120 and 160 active users, indicating an active user can absorb between 1.7 and 1.2 percent of the processor's capacity. The percent processor utilization values can be normalized to ARIPs for purposes of comparison across processors (See 'LARGE SYSTEM ARIPs "Approximate Relationships in Performance"', ZZ05-0386).

Main storage utilization in VM is not a very meaningful measurement. Therefore, we have focused our attention on user working set size and system paging rates. Normally, user working sets are in the 100-120 k byte (25-30 pages) range and are associated with moderate to low paging rates. Paging is a function of main storage access demand. As the number of concurrent users increases, the amount of storage available to each decreases until page faulting occurs with essentially every access. Page thrashing conditions associated with working set sizes of 70 k bytes have been observed.

The pagable portion of main storage, or the Dynamic Paging Area, constitutes approximately 12 of the 16 megabytes of the system's main storage configuration. The other 4 megabytes include the nucleus, internal trace table, and free storage area, with free storage requiring more than 3 megabytes.

I/O interrupt rates of up to 1250 per second have been observed. That translates into DASD I/O activity rates of 5-7 per active user per second, or 2-3 I/Os per second for user and system (non-paging) I/O and the remaining 3-4 per second for paging activity to drums. The translation of DASD I/O rates into percent capacity of the device is generally not straightforward. The capacity of a DASD to do I/O is a function of the device characteristics, the pathway configuration, data block size, and access methodology. In this environment, however, the performance characteristics of the three types of DASD (IBM 3380 disk, IBM 2305 drum, and STC 4305 solid state drum) are the primary variables. The I/O pathway configuration is relatively stable, the data block size is fixed by function and is generally 2 k bytes for user data and 4 k bytes for paging, and the access methods are constant.

Plots of device utilization as a function of I/O activity and I/O activity queued as a function of device utilization indicate device capacities to do I/O. Comparison of values for different I/O and processor configurations indicate the effect of those components on I/O device capacity.

Those plots confirm the 30-35 percent device busy rule of thumb for 3380s and indicate an 'average' capacity of 12-15 I/Os per second, device service times of 20-25 milliseconds, and I/O activity queued of less than or equal to 10 percent. ('Average' values are indicated for a 9-hour day with peak period values of 2 to 2.5 times the 'average'.)

For 2305 drums, used as paging devices with three exposures, 10 percent I/O queued occurs around 45 percent device busy, indicating an 'average' capacity of 25-35 I/Os per second depending on the system configuration. Device service times of 10-20 milliseconds are observed. The highest I/O capacity values are available from devices on dedicated pathways on 3033 processors, with lower values for devices on 3081 processors or contended pathways.

For the solid state drums, used as paging devices with three exposures, projections of logical device busy to 45-50 percent indicates an 'average' I/O capacity of greater than 60 per second. Device service times on data streaming channels at approximately 30 I/Os per second are between 3 and 5 milliseconds.

Capacity Observations

A gross level performance analysis of the VM system indicated an imbalance between CPU and main storage capacity. Between 120 and 160 active users can be supported by a 190-200 percent utilized processor, but only 120 active users with working set sizes of 100 k bytes can be supported in the 12 megabytes of dynamically pagable storage. The support of 160 active users (40 beyond the 'capacity' of main storage) is accomplished through the high performance characteristics of the paging subsystem. However, at the high paging rates associated with 160 active users, the processor resources are increasingly consumed by CP paging overhead, thereby resulting in decreasing amounts of virtual (problem program) time to accomplish user work.

The non-paging DASD configuration to support 120 active users is estimated to be between 24 and 30 arms (3 I/Os per active user per second * 120 active users / 15 or 12 I/Os per second per device). For our environment with average user minidisks in the 3-5 cylinder range and the low ratio of active to total users, the 3380's I/O and storage capacity are reasonably balanced for user minidisk storage. However, the storage capacity of the 3380 is excessive relative to its I/O capacity for higher density of data reference system functions such as temporary and S and Y disk accesses. Multiple, partially allocated disks are dedicated to those functions to obtain the desired performance. Forty-four 3380 arms are configured on the 3081 system to support the I/O load and accommodate the I/O access skew across devices.

The paging DASD configuration to support 120 active user at 4 pages per user per second would be 20 devices at 25 pages per second, 12 devices at 40 pages per second, and 8 devices at 60 pages per second. Our experience shows that IBM 2305 drums have the I/O capacity of 25 to 35 I/Os per second and that solid state drums are projected to have an I/O capacity of 60 I/Os per second.

The I/O pathway configuration to support the nonpaging I/O load is estimated to be approximately 6 pathways (channel/control unit) - 360 I/Os * .004 sec per I/O / 25 percent pathway component busy, or 6 to 8 arms per path. Ten channel pathways (5 symmetrical pairs) are configured to support the current I/O load requirement as well as to provide for future growth.

For the paging I/O load, approximately 1 pathway per logical device is indicated for rotating drums and 1 pathway per 2 logical devices for solid state drums. However, both the 2305 and the 4305 are configured with multiple logical devices per controller, 2 on the 2305 and either 2 or 4 on the 4305. The current configuration includes 12 logical solid state drums accessed via 6 channel pathways. Previous 2305 drum configurations had 8 logical drums accessed via 8 channels.

The capacity of the channel/control unit path to do I/O is empirically determined by comparison of device service times with pathway component utilization. Initially, maximum tolerable device service time values are specified. Then, by adding workload or devices, I/O pathway loading is increased until the maximum specified device service times are reached. Device service time increases concomitantly with pathway component utilization because of increasingly frequent missed device reconnects. The resultant pathway loading values (I/O rates) are considered to be capacity values. Preliminary results indicate that 'average' pathway capacity is approximately 50 I/Os per second for rotating devices and 100 I/Os per second for solid state devices, the difference being the penalty of a full rotation for missed reconnects.

Future Work

The near term future requires the development of an I/O configuration that will support the growing CMS intensive workload on a 3081 K with 32 megabytes of main storage and continue to provide the same or a higher performance environment.

The longer term project is the characterization of the VM user's terminal utilization. The intent is to determine the time a user is logged on the system, length of session, number of interactions per session, and user 'think' time. We anticipate that this information will provide insight into the question of the number of users a terminal can support and, correspondingly, into system resource requirements per system user or per terminal. This information will also assist in the exploration of the relationship of system performance and user productivity in this VM/CMS environment.

A still longer term, lower priority, task is the characterization of the functions of a computer performance analyst. This information may be instrumental in the establishment of an effective process of identification and development of performance analysts to satisfy the growing demand.

VM PERFORMANCE ANALYSIS AND CAPACITY PLANNING WITH EMPHASIS ON DASD

KEITH E. SILLIMAN

IBM CORPORATION
FEDERAL SYSTEMS DIVISION
18100 FREDERICK PIKE
GAITHERSBURG, MARYLAND 20879
8-372-7132
301-840-7132

INTRODUCTION

VM PERFORMANCE ANALYSIS & CAPACITY PLANNING METHODOLOGY
CPU, MAIN STORAGE, & DASD MEASUREMENT RESULTS
CAPACITY OBSERVATIONS — CPU, MAIN STORAGE, & DASD
FUTURE DIRECTIONS

Sanitized Copy Approved for Release 2010/06/07 : CIA-RDP85-00142R000300200003-6

INTRODUCTION

LARGE SCALE, HIGH PERFORMANCE VM/CMS SYSTEM
3033MP/3081 D
TRIVIAL RESPONSE THRESHOLD — 0.3 SECONDS

Sanitized Copy Approved for Release 2010/06/07 : CIA-RDP85-00142R000300200003-6

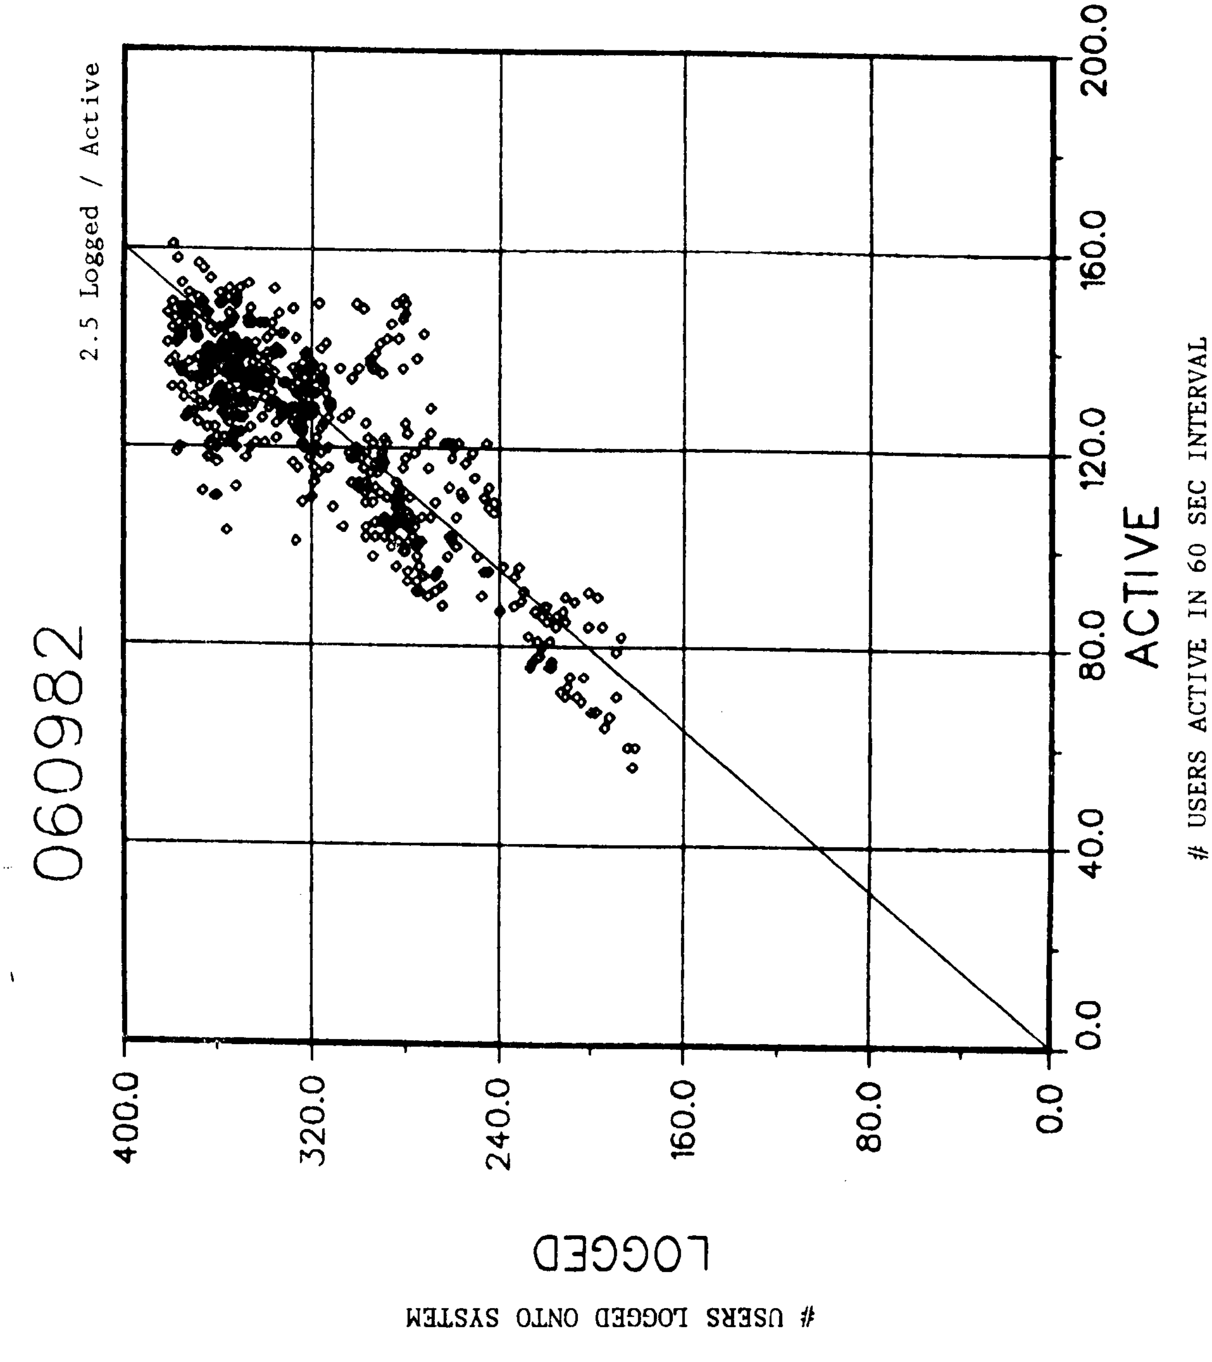

Sanitized Copy Approved for Release 2010/06/07 : CIA-RDP85-00142R000300200003-6

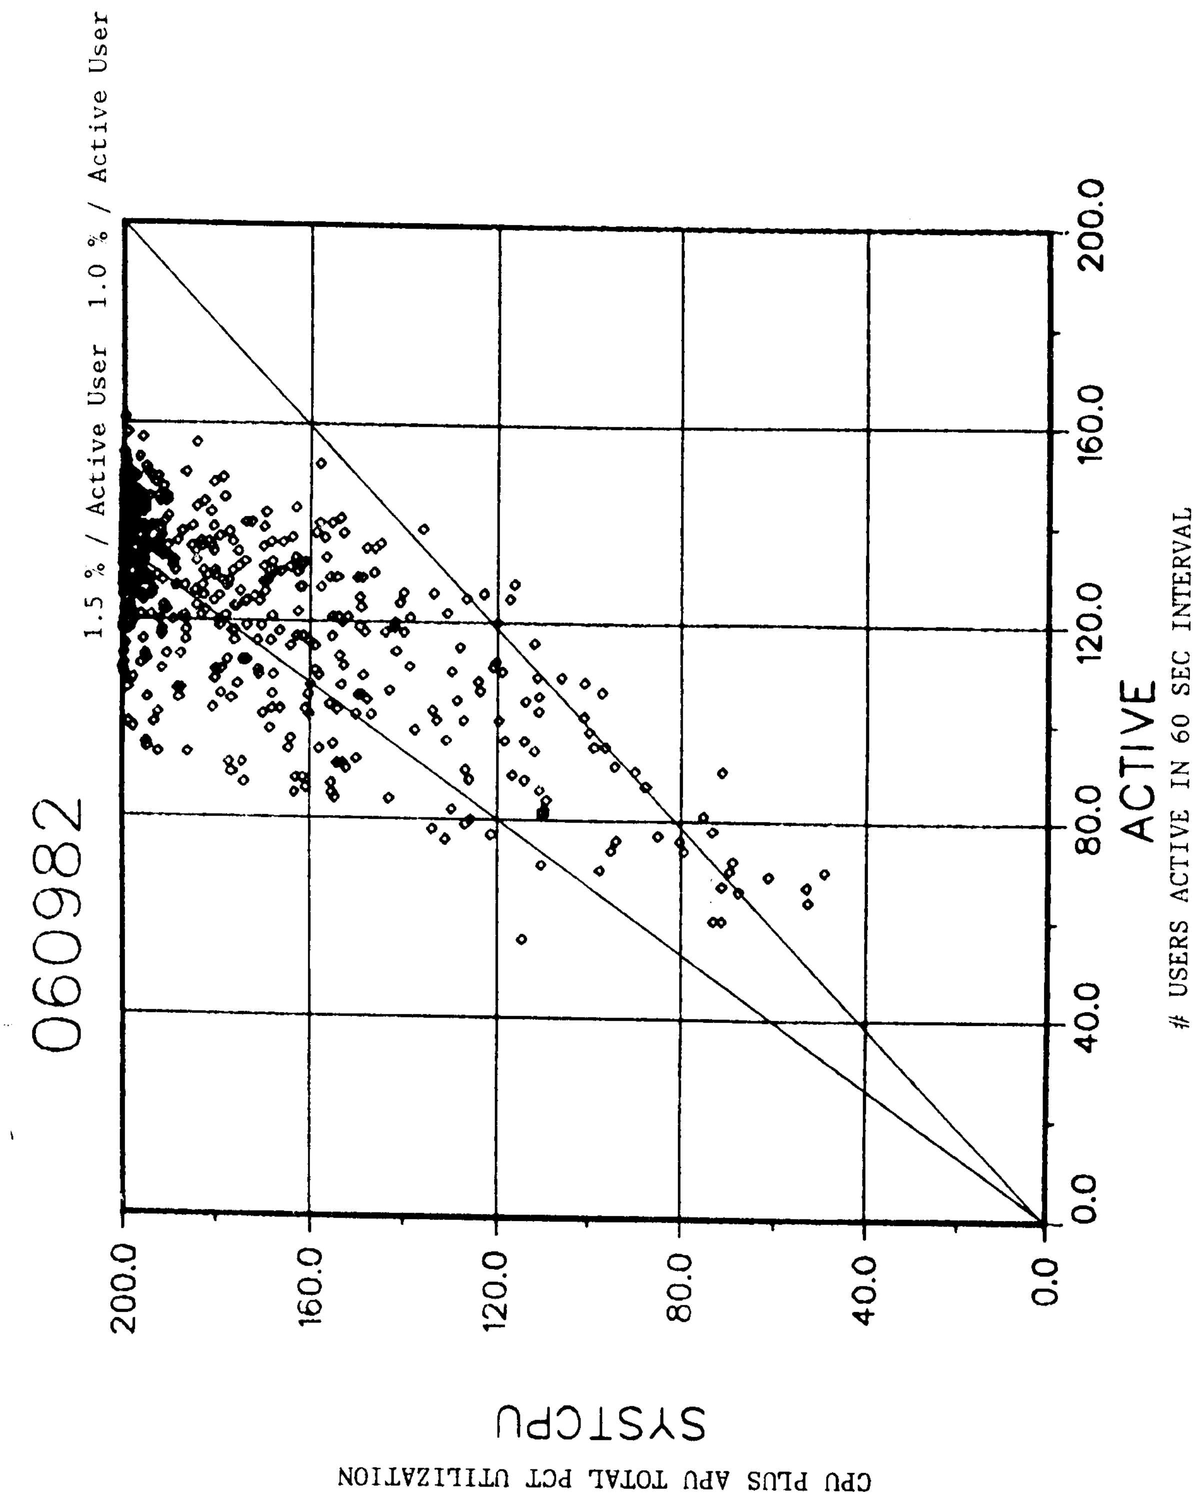
060982
1.5 % / Active User 1.0 % / Active User
ACTIVE
USERS ACTIVE IN 60 SEC INTERVAL
0.0
40.0
80.0
120.0
160.0
200.0
SYSTCPU
CPU PLUS APU TOTAL PCT UTILIZATION

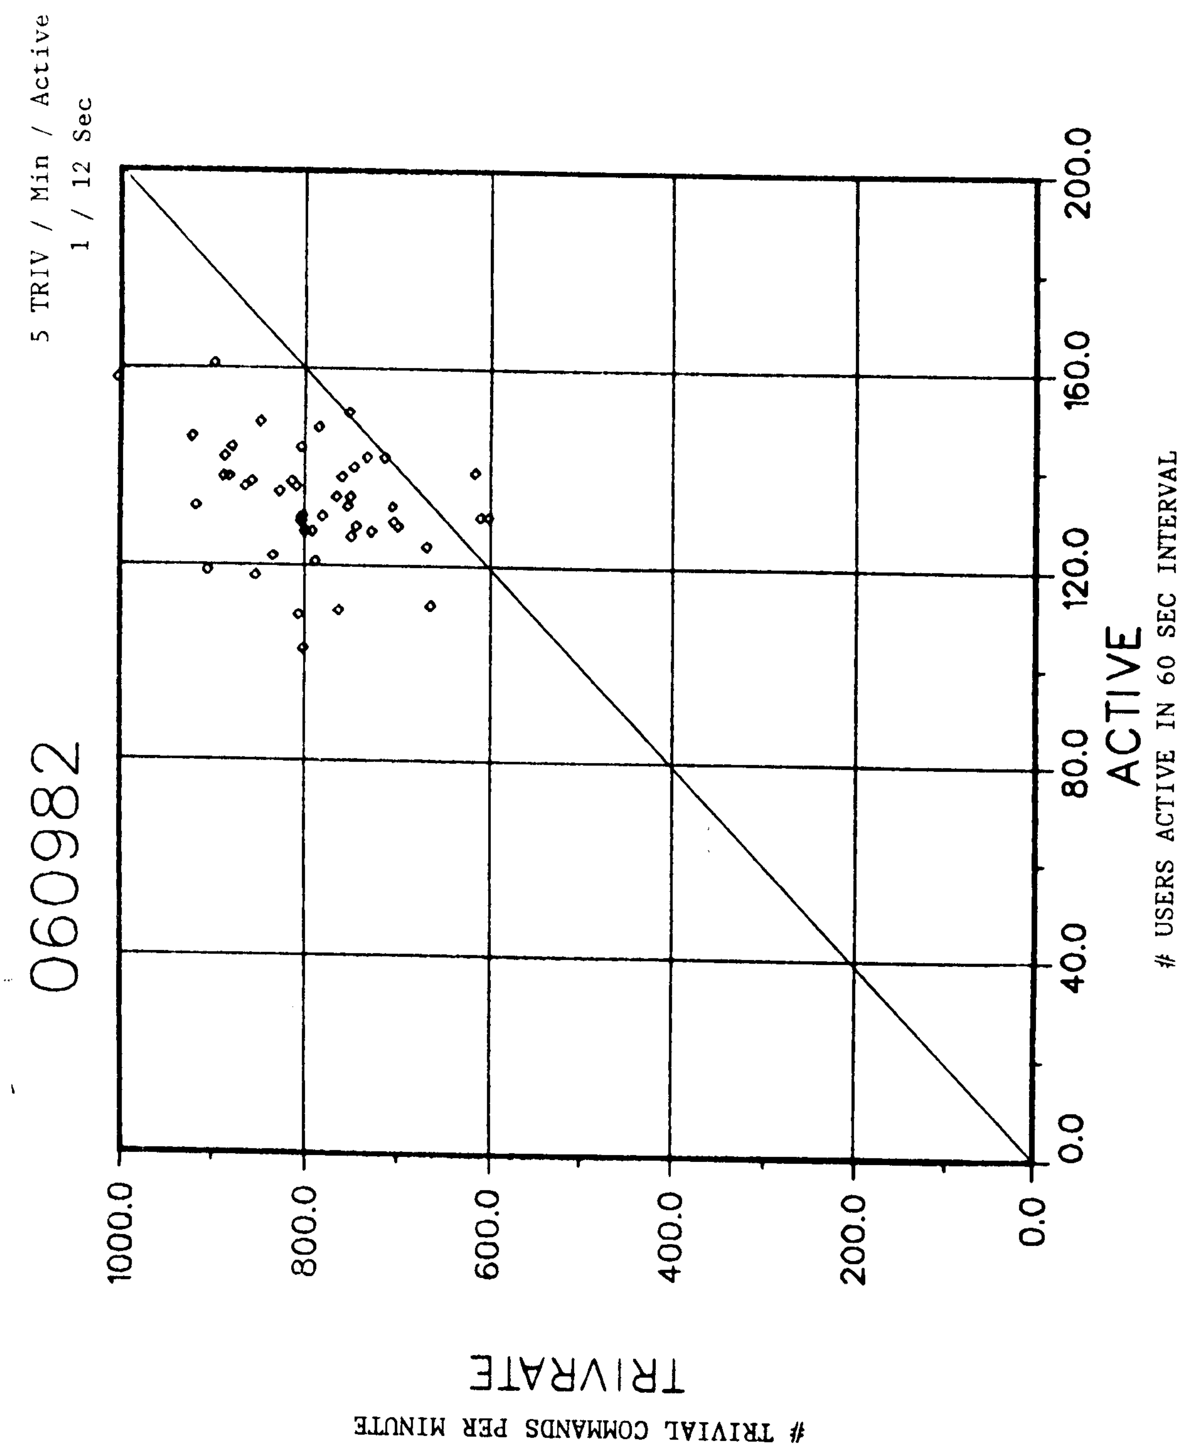
060982
5 TRIV / Min / Active
1 / 12 Sec
ACTIVE
USERS ACTIVE IN 60 SEC INTERVAL
TRIVRATE
TRIVIAL COMMANDS PER MINUTE
0.0
40.0
80.0
120.0
160.0
200.0
0.0
200.0
400.0
600.0
800.0
1000.0

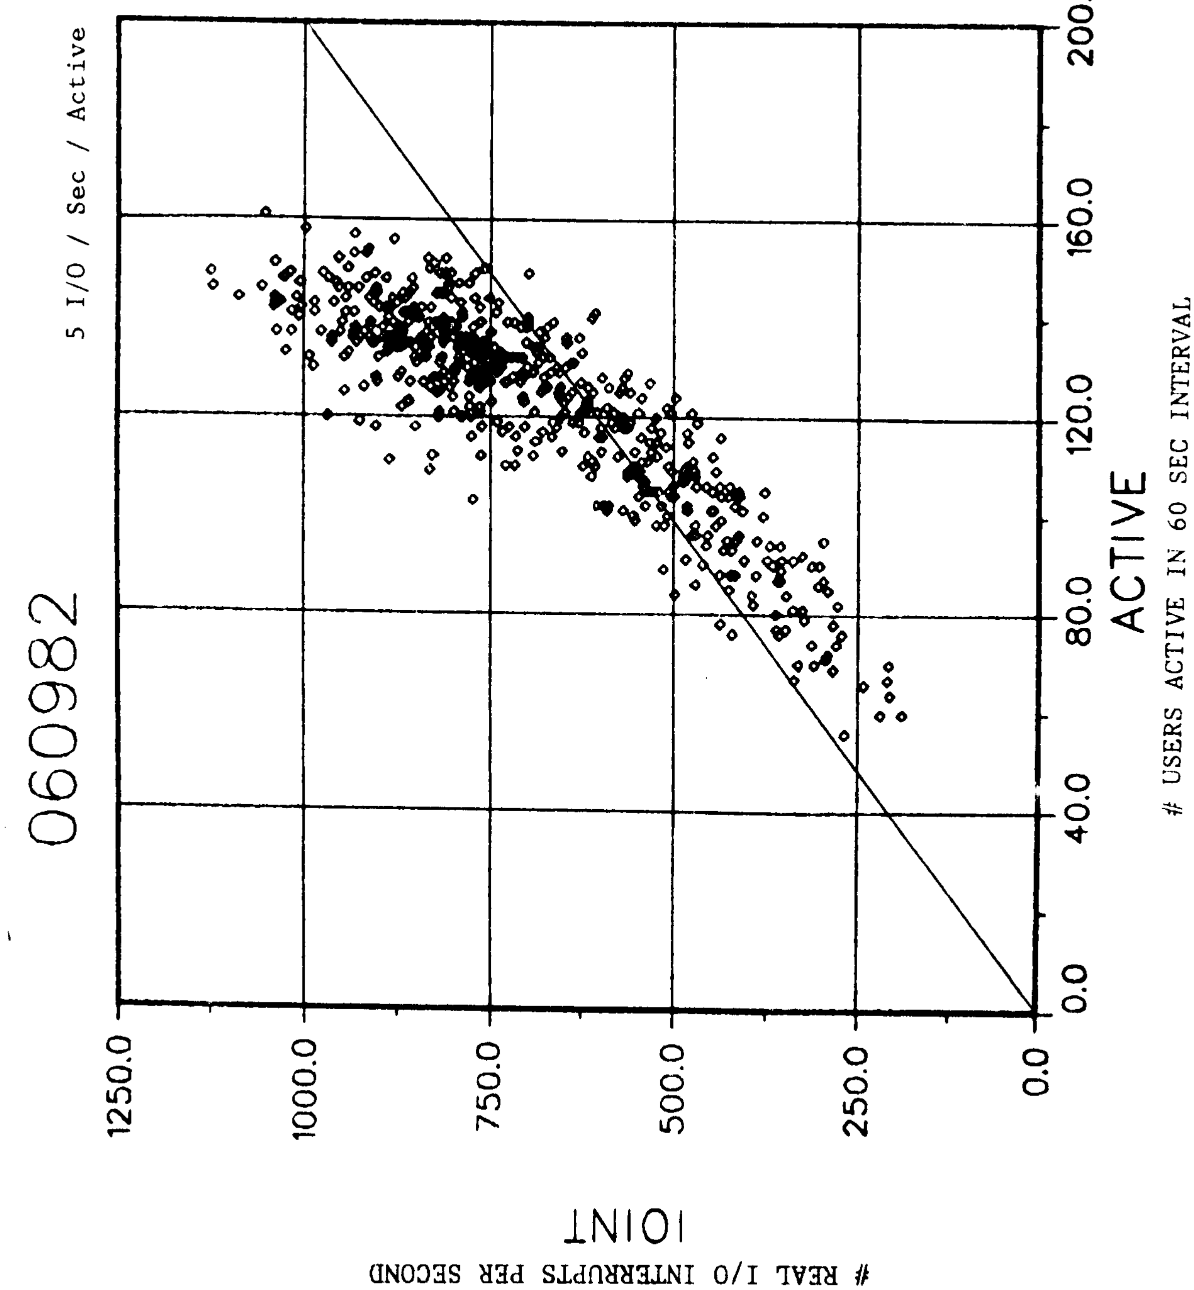

Sanitized Copy Approved for Release 2010/06/07 : CIA-RDP85-00142R000300200003-6

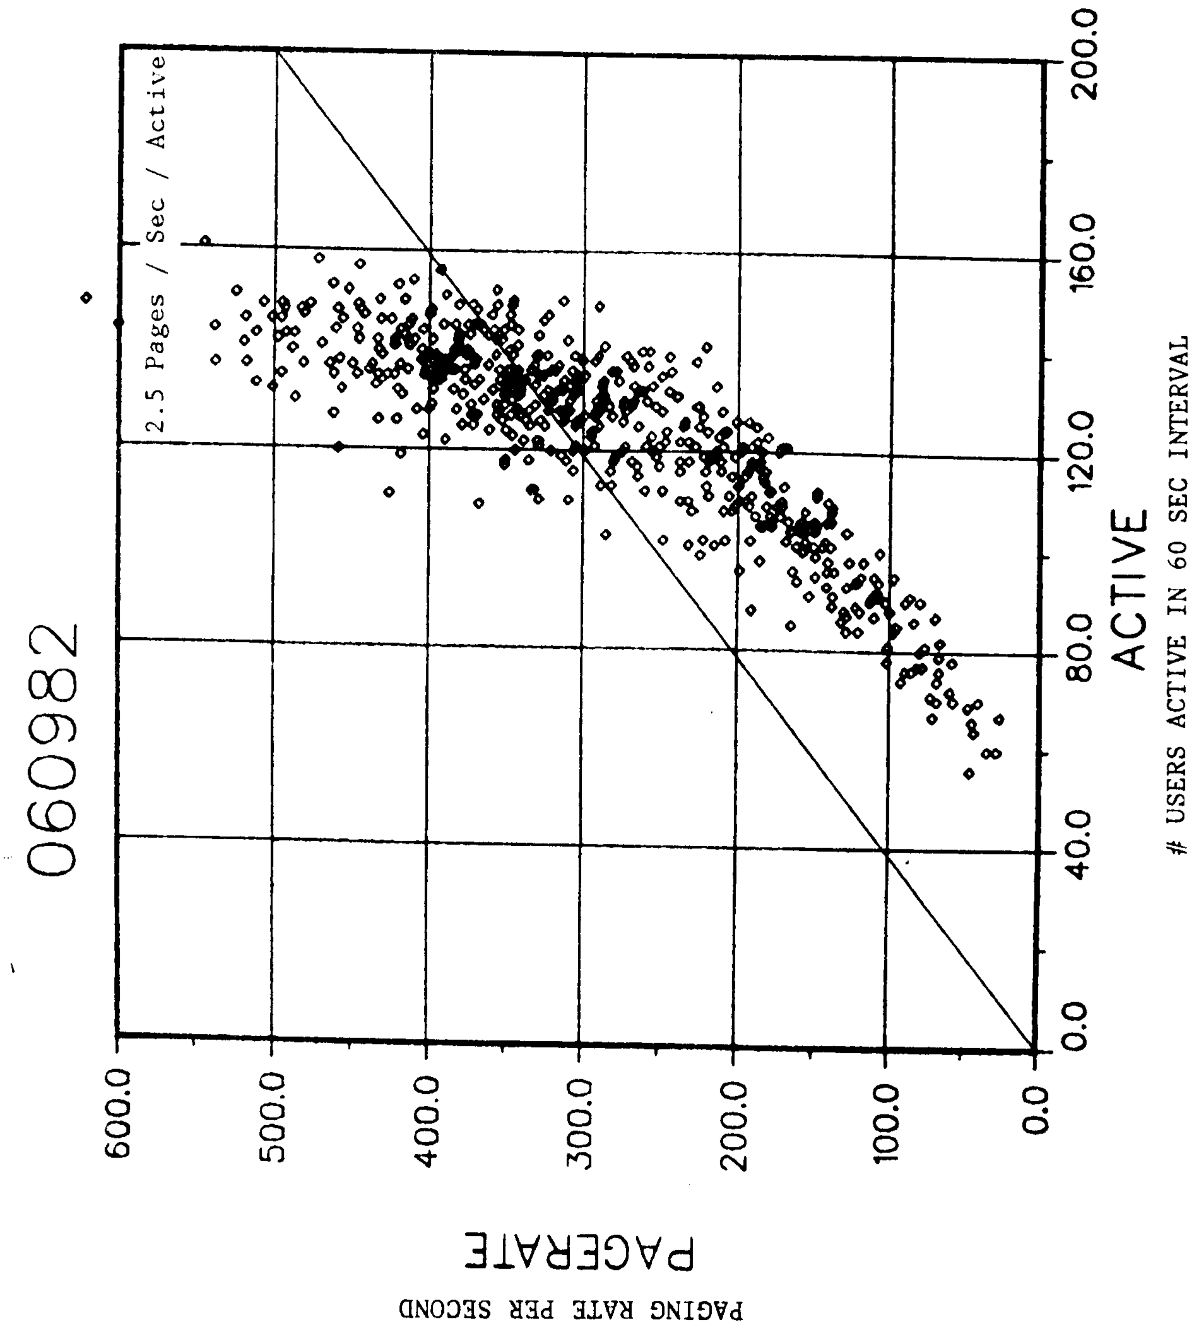
060982
2.5 Pages / Sec / Active
0.0
40.0
80.0
120.0
160.0
200.0
ACTIVE
USERS ACTIVE IN 60 SEC INTERVAL
600.0
500.0
400.0
300.0
200.0
100.0
0.0
PAGERATE
PAGING RATE PER SECOND

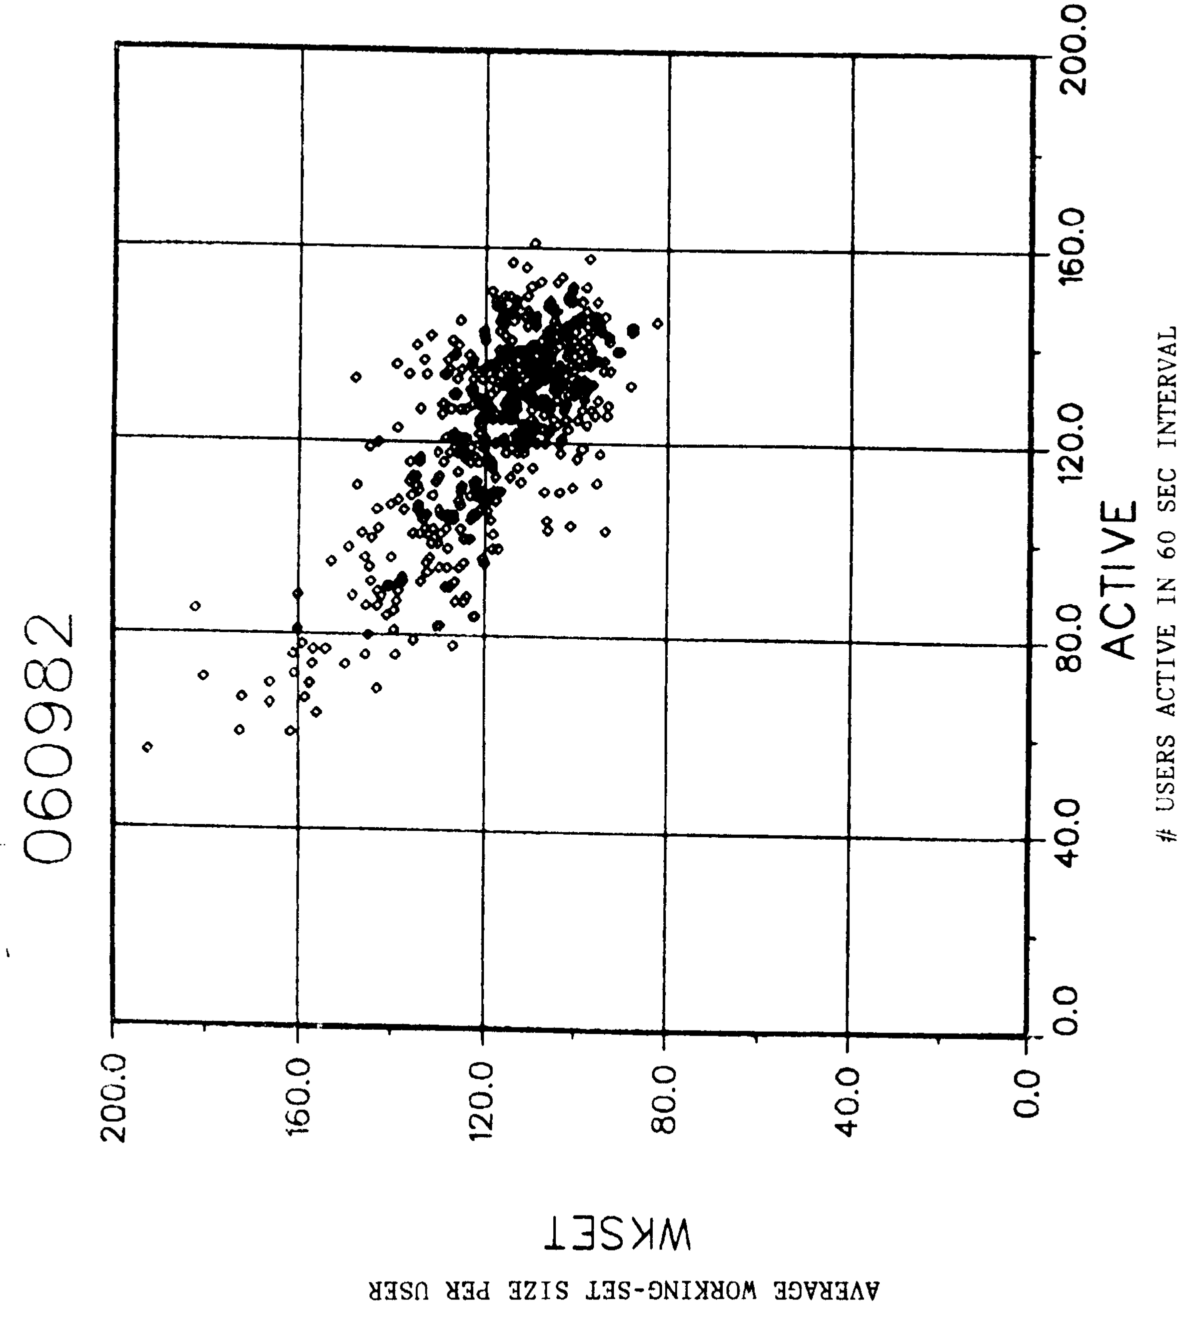

060982
ACTIVE
USERS ACTIVE IN 60 SEC INTERVAL
0.0
40.0
80.0
120.0
160.0
200.0
WKSET
AVERAGE WORKING-SET SIZE PER USER

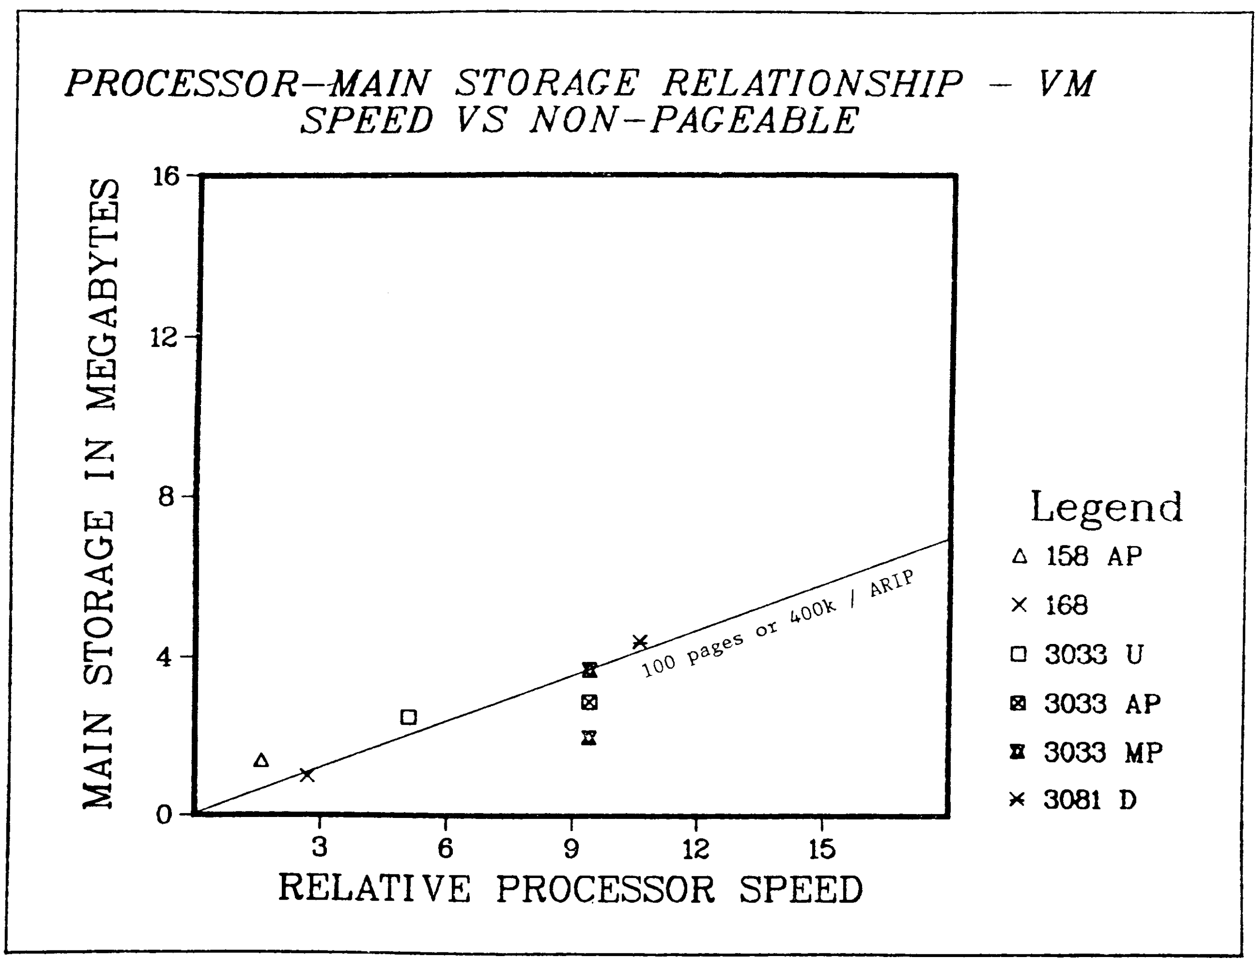
PROCESSOR-MAIN STORAGE RELATIONSHIP - VM
SPEED VS NON-PAGEABLE
MAIN STORAGE IN MEGABYTES
16
12
8
4
0
3
6
9
12
15
RELATIVE PROCESSOR SPEED
100 pages or 400k / ARIP
Legend
158 AP
168
3033 U
3033 AP
3033 MP
3081 D

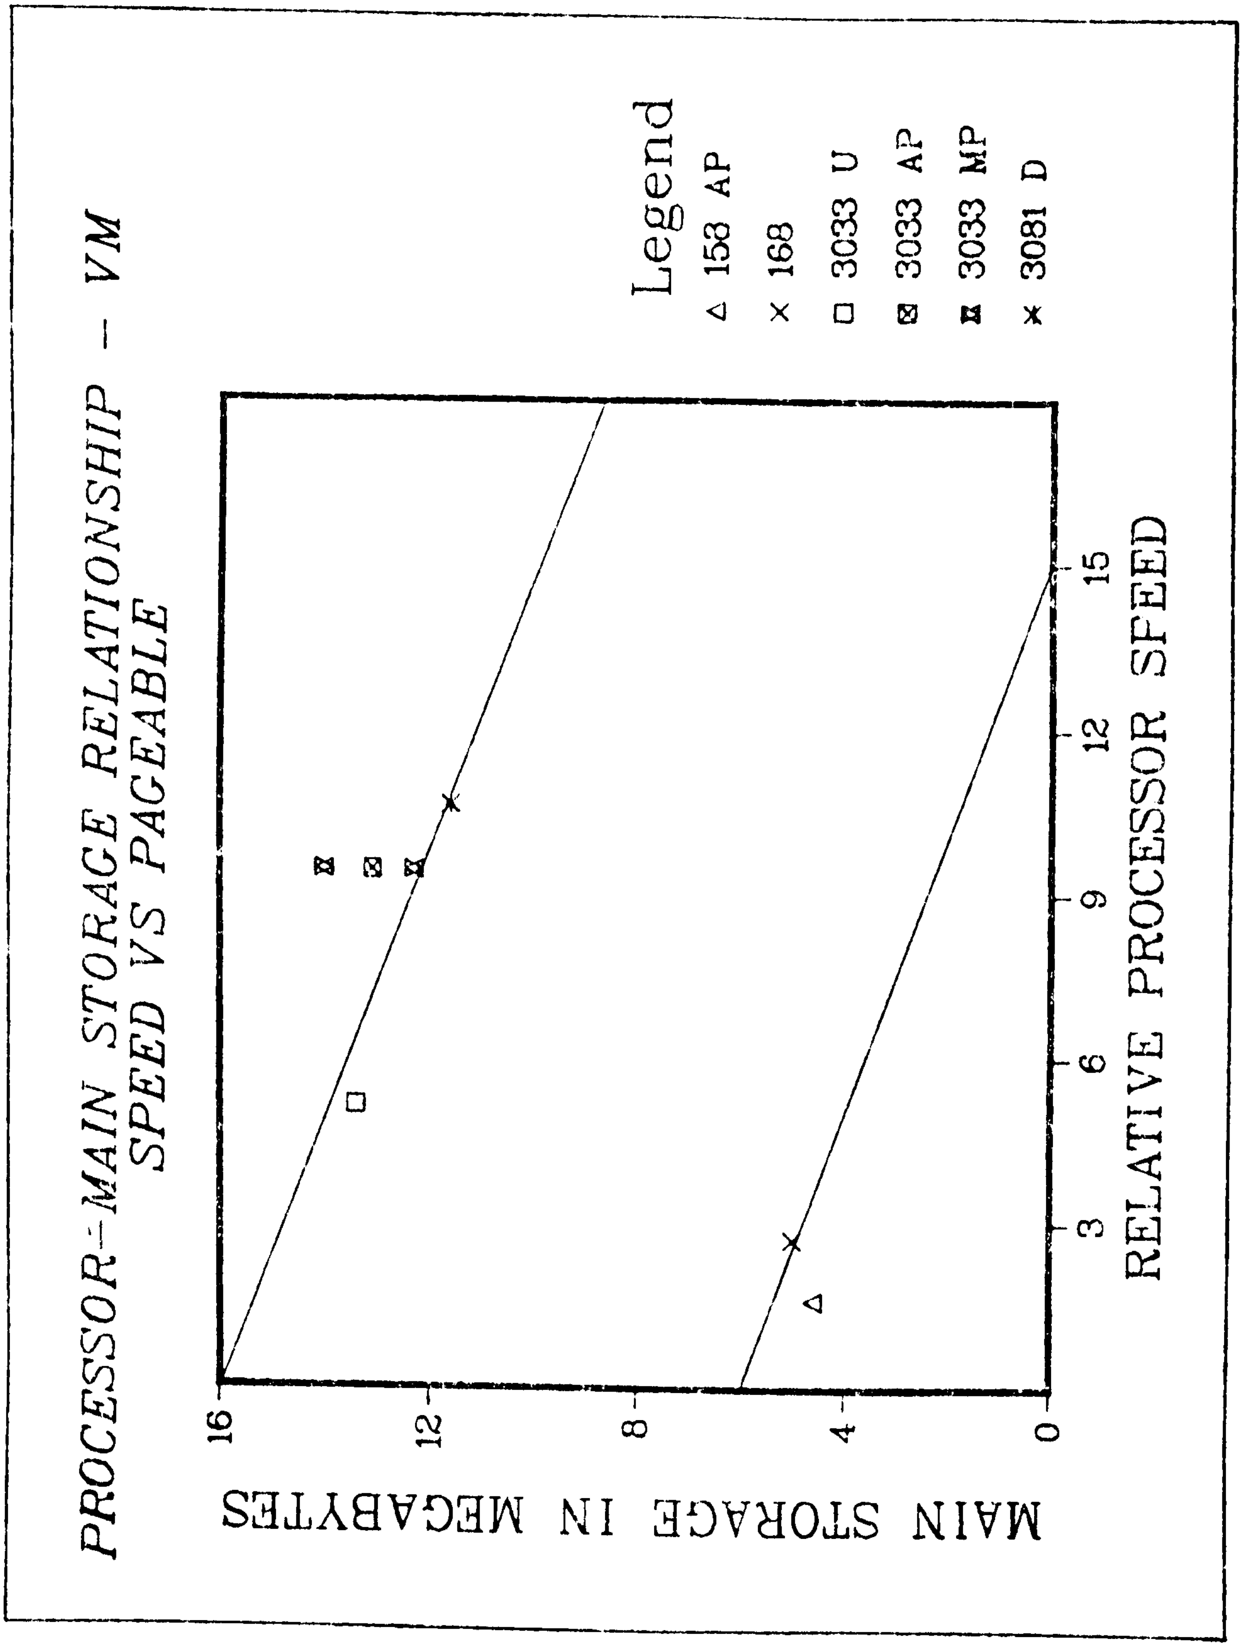
PROCESSOR-MAIN STORAGE RELATIONSHIP - VM
SPEED VS PAGEABLE
MAIN STORAGE IN MEGABYTES
RELATIVE PROCESSOR SPEED
0
4
8
12
16
3
6
9
12
15
Legend
△ 153 AP
× 168
□ 3033 U
⊠ 3033 AP
⌧ 3033 MP
✱ 3081 D

PROCESSOR–MAIN STORAGE RELATIONSHIP — VM
SPEED VS PAGEABLE

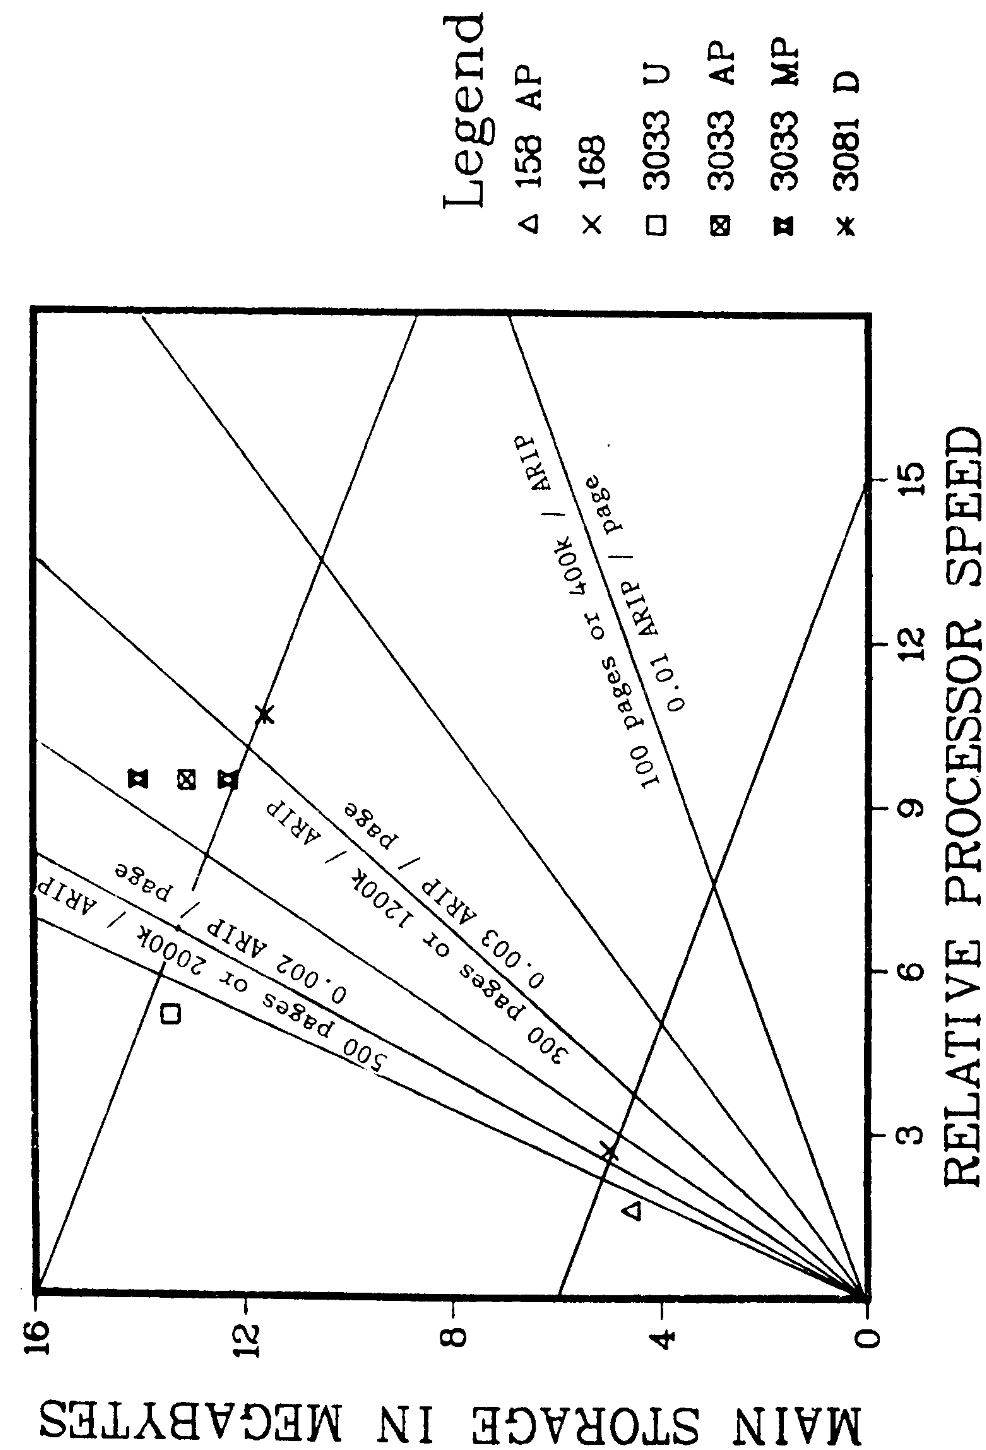

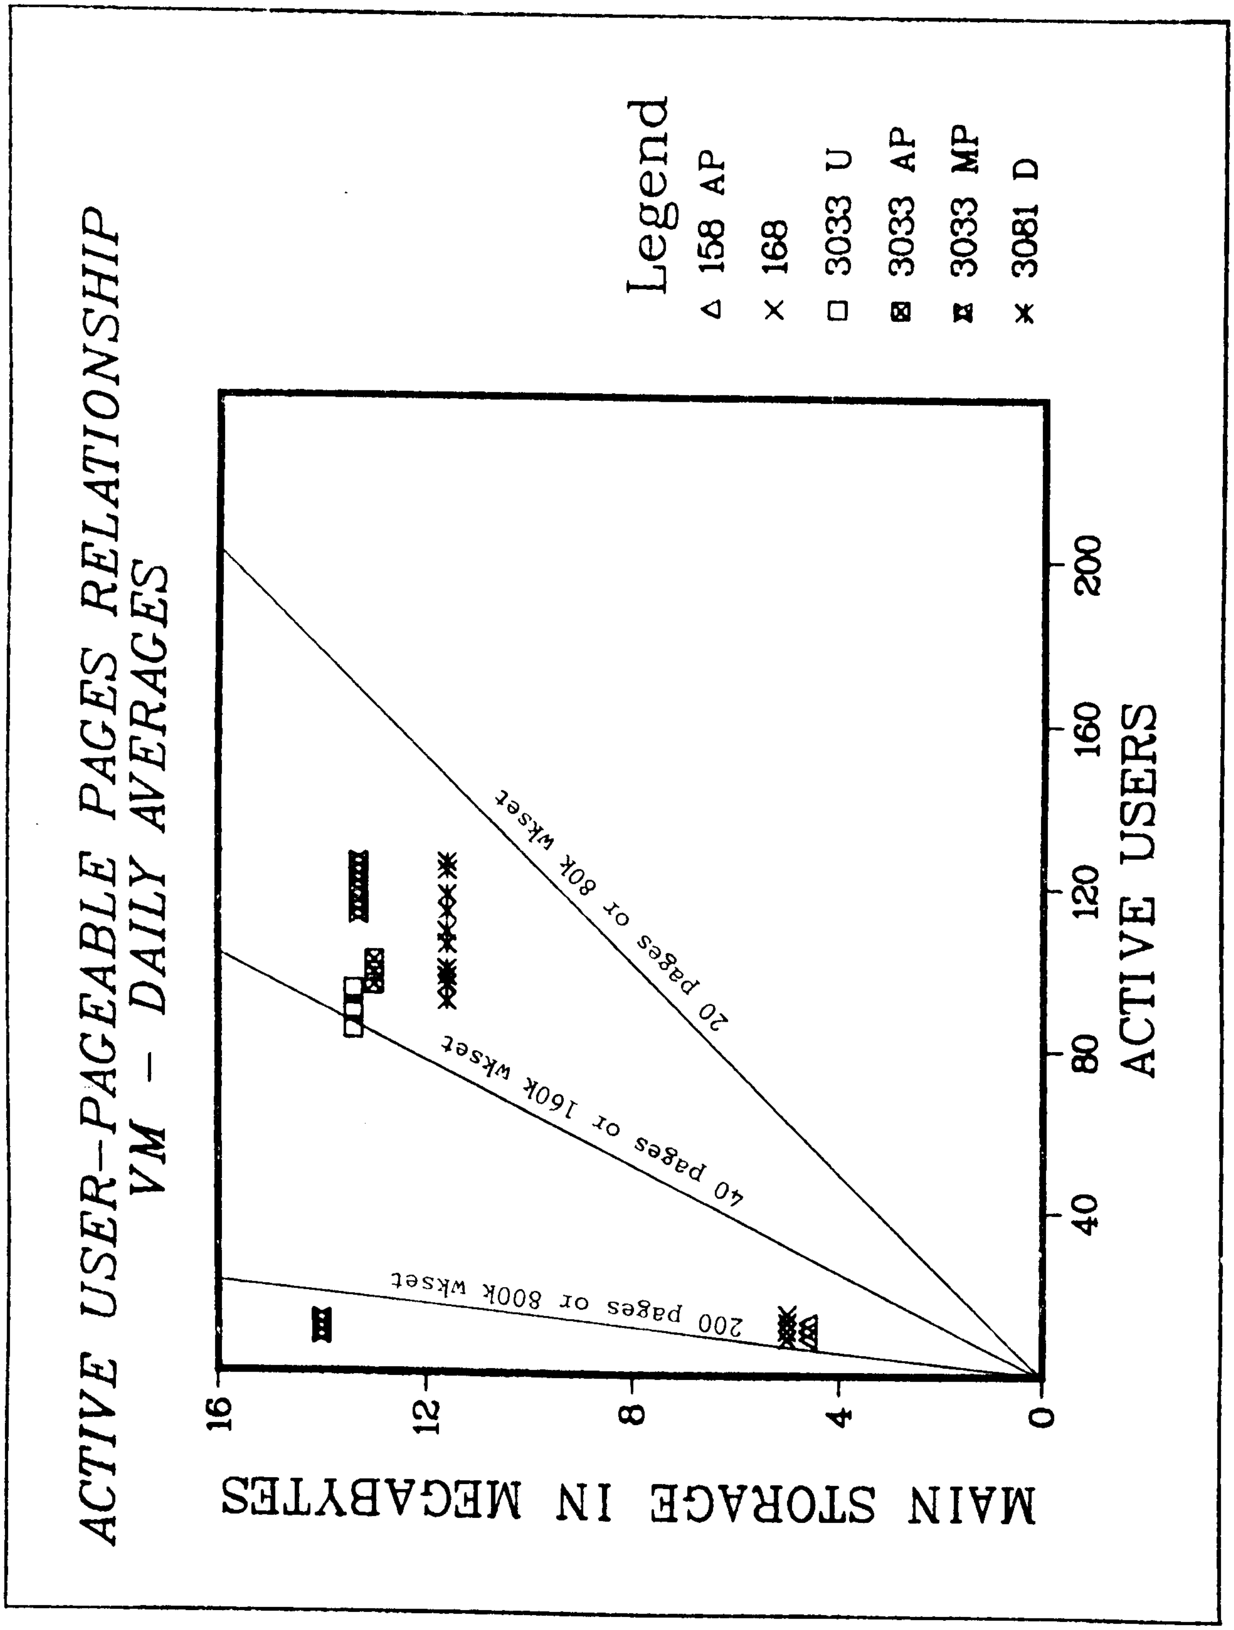
ACTIVE USER-PAGEABLE PAGES RELATIONSHIP
VM - DAILY AVERAGES
Legend
△ 158 AP
× 168
□ 3033 U
⊠ 3033 AP
3033 MP
✱ 3081 D
20 pages or 80k wkset
40 pages or 160k wkset
200 pages or 800k wkset
MAIN STORAGE IN MEGABYTES
0
4
8
12
16
ACTIVE USERS
40
80
120
160
200

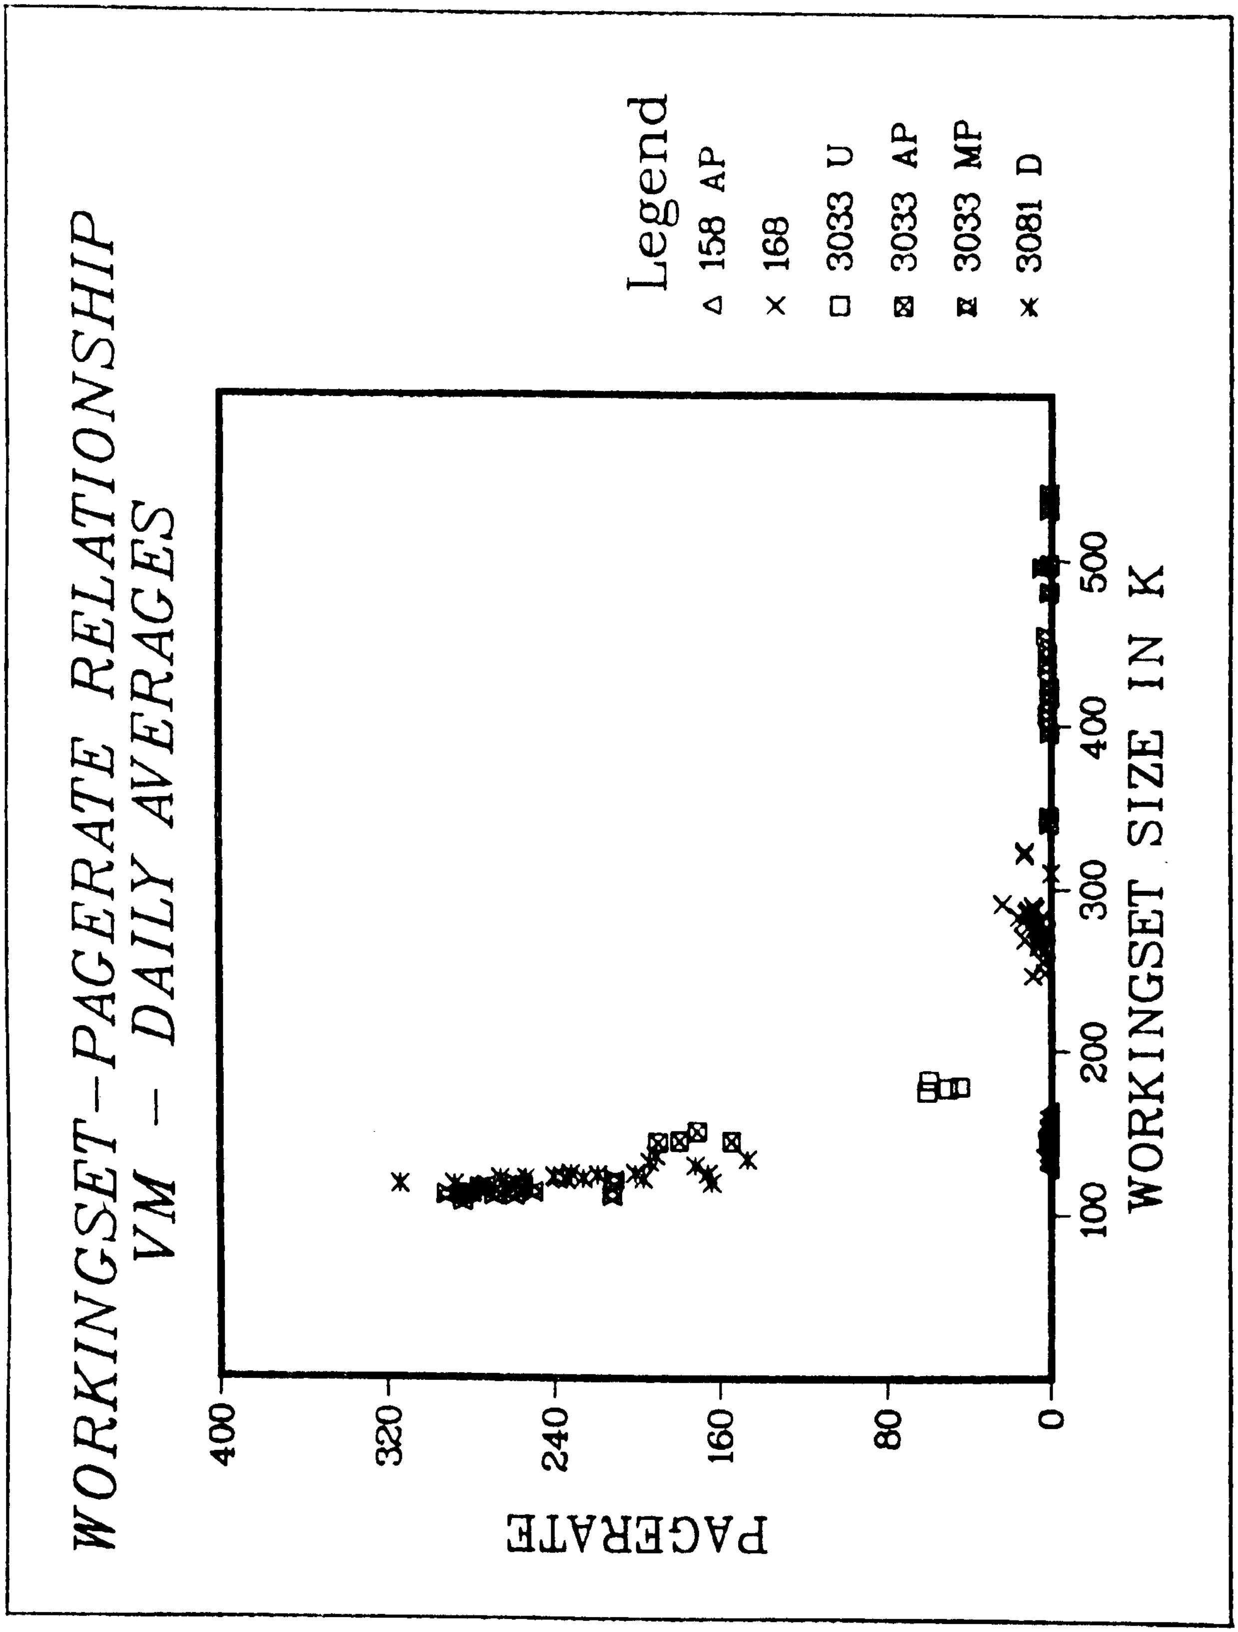
WORKINGSET-PAGERATE RELATIONSHIP
VM - DAILY AVERAGES
Legend
158 AP
168
3033 U
3033 AP
3033 MP
3081 D
PAGERATE
0
80
160
240
320
400
WORKINGSET SIZE IN K
100
200
300
400
500

MAIN STORAGE AND PAGE RATES IN VM

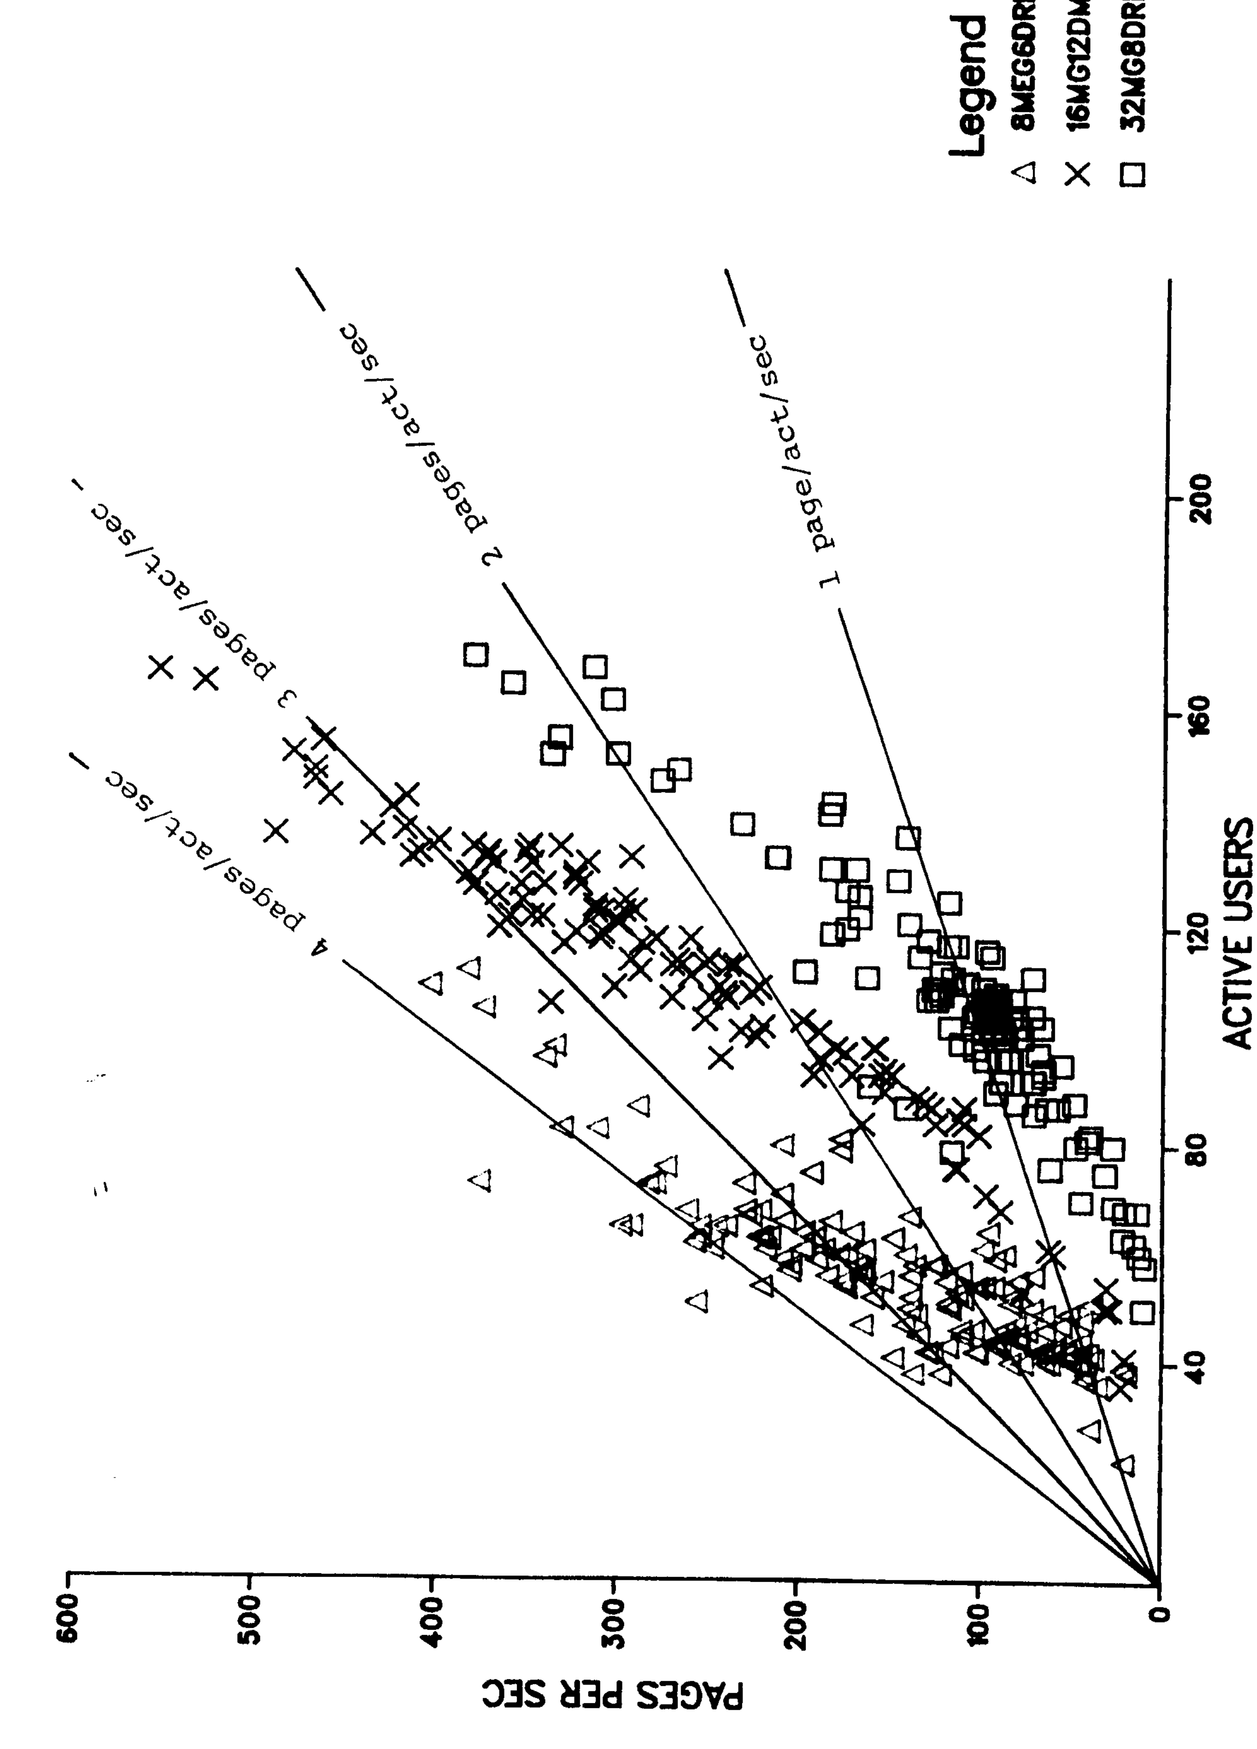

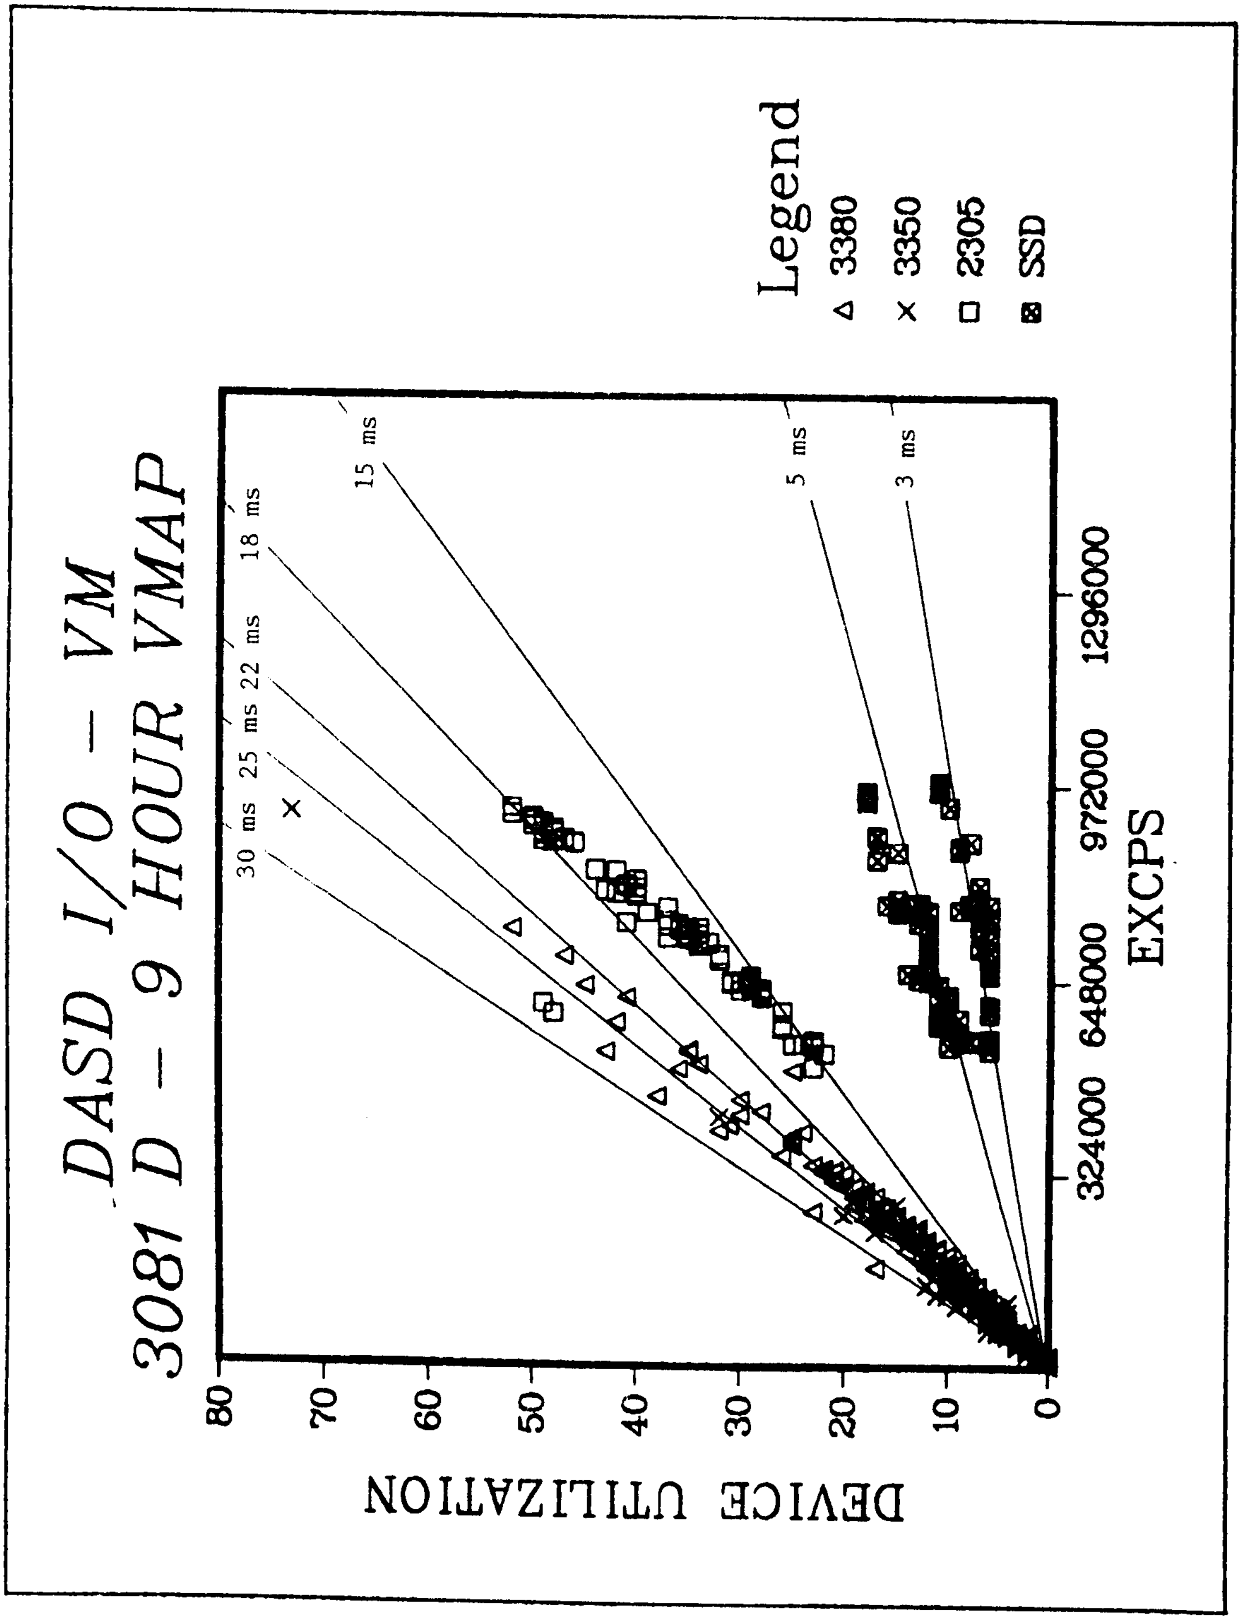
DASD I/O – VM
3081 D – 9 HOUR VMAP
Legend
3380
3350
2305
SSD
30 ms
25 ms
22 ms
18 ms
15 ms
5 ms
3 ms
DEVICE UTILIZATION
0
10
20
30
40
50
60
70
80
EXCPS
324000
648000
972000
1296000

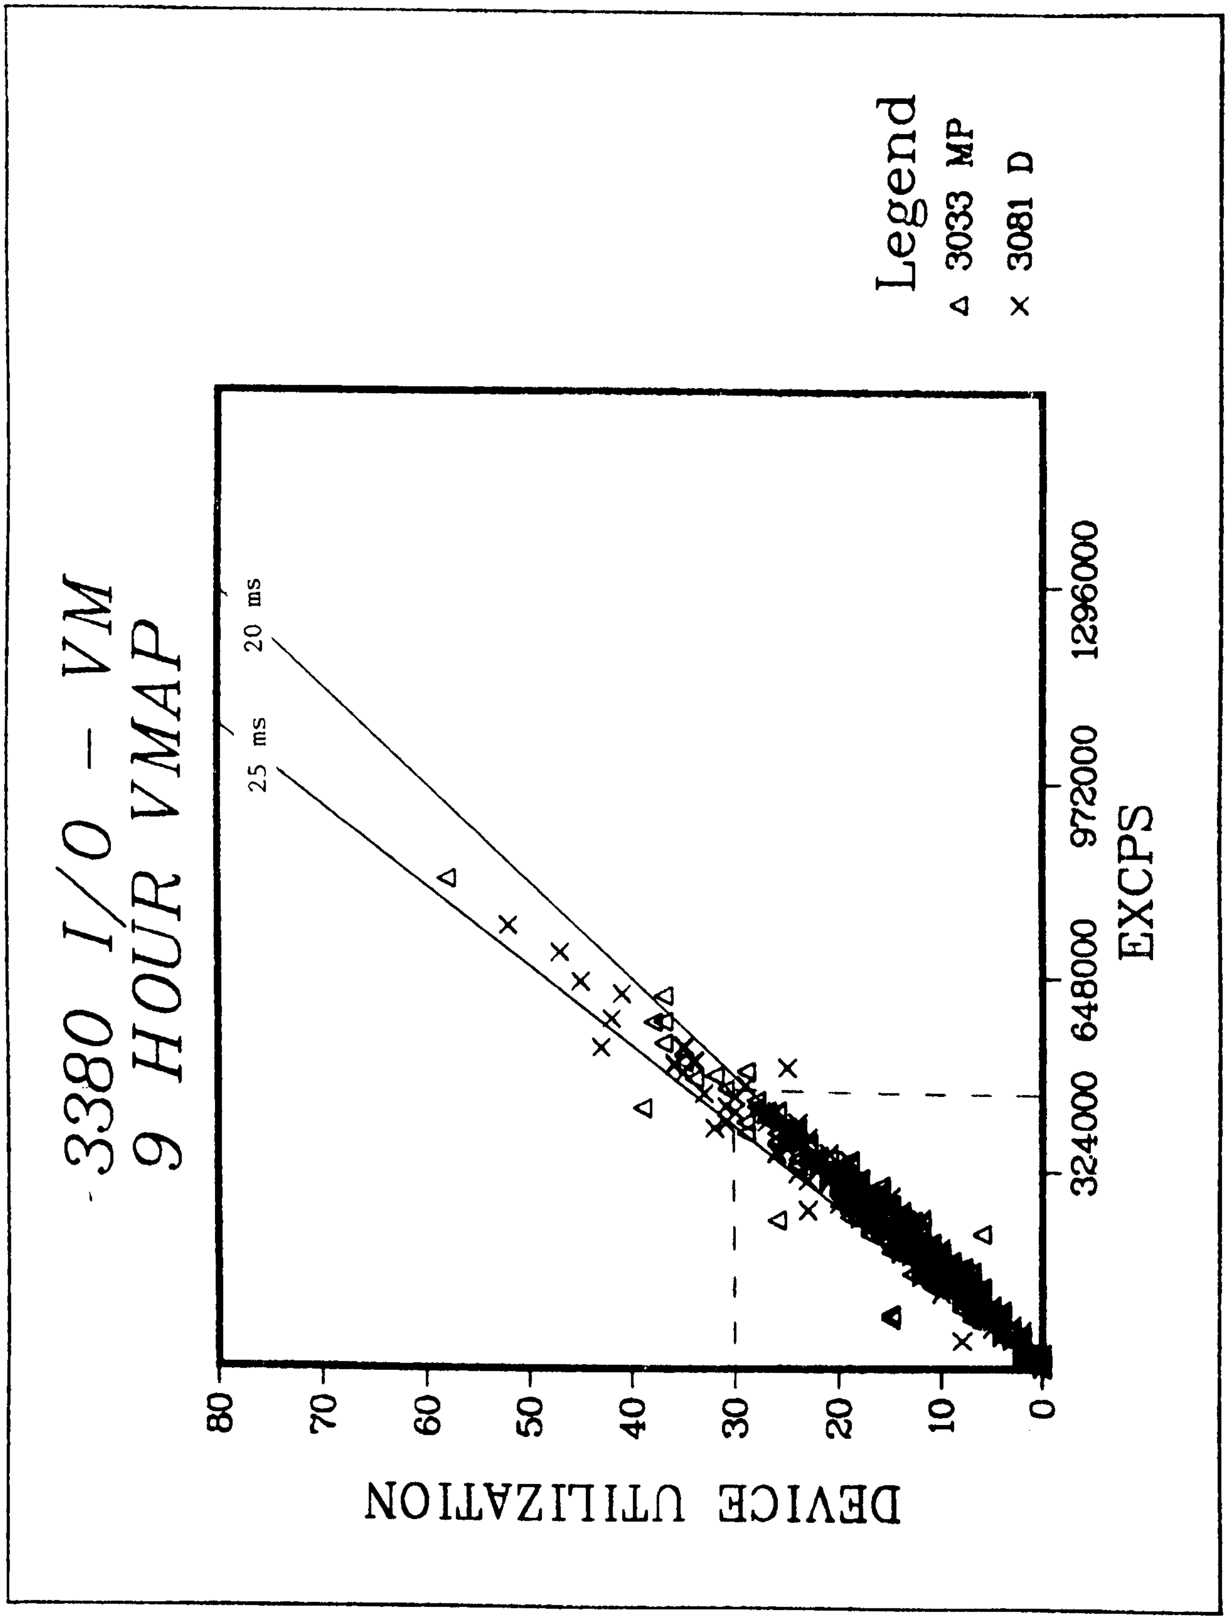

Sanitized Copy Approved for Release 2010/06/07 : CIA-RDP85-00142R000300200003-6

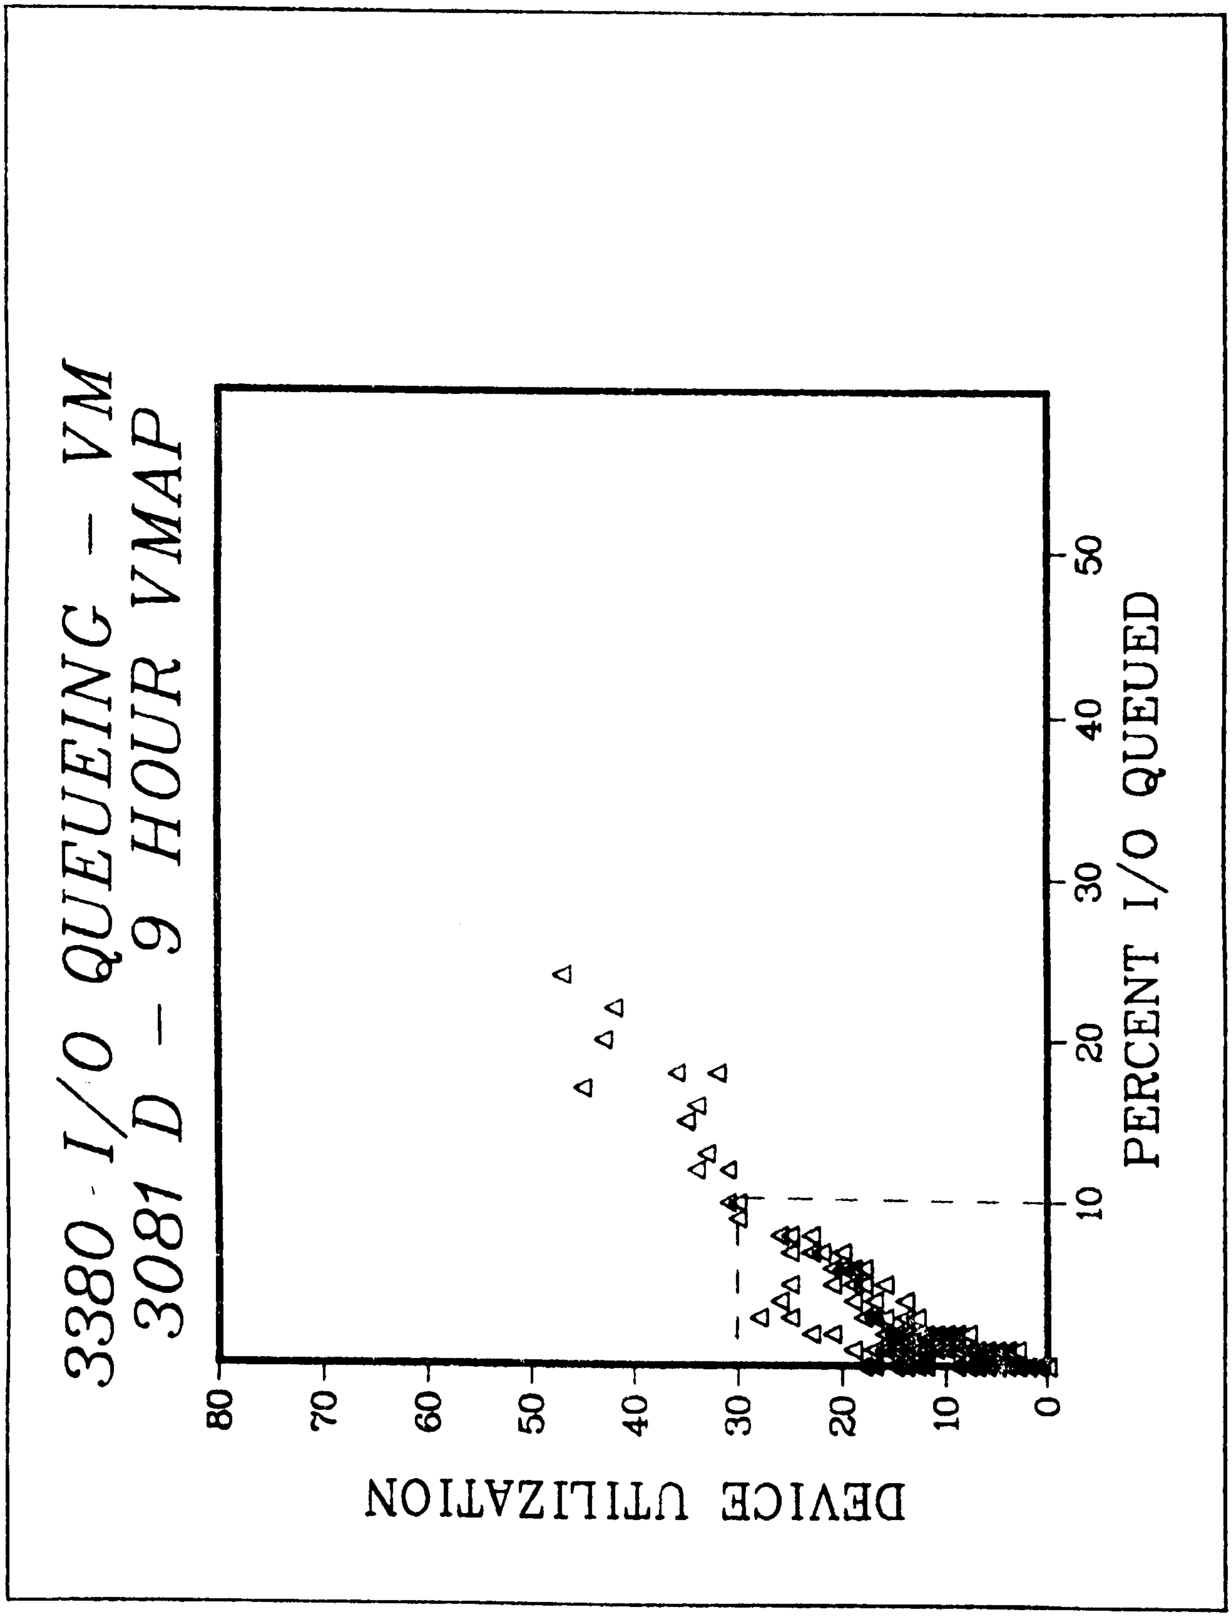
3380 I/O QUEUEING — VM
3081 D — 9 HOUR VMAP
DEVICE UTILIZATION
0
10
20
30
40
50
60
70
80
PERCENT I/O QUEUED
10
20
30
40
50

COMBINED I/O ACTIVITY, VM1 & VM2
Thursday June 30, 1983

DEV ADD	VOL ID	VM1 I/O	VM2 I/O	TOTAL I/O	I/O SEC	VM1 %DU	VM2 %DU	TOT %DU	VM1 %QD	VM2 %QD	SERVICE TIME ms
920	VMBKP1	233348		233348	7	18		18	1		22
921	V2U802	41075	25421	66496	2	2	2	4	0	0	17
B22	VMU123	14055	7827	21882	1	0	0	0	0	0	
B23	VMU127	16939	66879	83818	3	1	4	5	0	0	15
924	VMU101	148557	36158	184715	6	10	3	13	1	0	17
925	VMU129	56079	56066	112145	3	4	4	8	1	1	20
B26	VMU103	73502	41496	114998	4	5	2	7	1	0	16
B27	V1AIM1	124119		124119	4	6		6	0		15
928	V1U106	10063	27532	37595	1	0	2	2	0	0	24
929	V1U107	143505	33978	177483	5	11	3	14	1	1	22
B2A	V1U108	66045	39869	105914	3	4	3	7	1	1	19
B2B	V1U109	57998	73216	131214	4	4	6	10	0	0	23
92C	V1IPL1	321146		321146	10	24		24	6		23
92D	V2U807	25478	72509	97987	3	2	6	8	1	1	19
B2E	V1IPL2	20956		20956	1	2		2	1		2
B2F	VMTRL1	57534	753018	810552	25	4	43	47	1	12	16
x2x		1410399	1233969	2644368	82	97	78	175			22
92x		979251	251664	1230915	38	71	20	91			24
B2x		431148	982305	1413453	44	26	58	84			19
940	VMU121	16092	42995	59087	2	0	2	2	0	0	7
941	V1BAL1	224456		224456	7	14		14	1		20
B42	VMU124	31083	51360	82443	3	2	2	4	0	0	11
B43	VMU128	90530	15359	105889	3	5	2	7	0	0	15
944	VMU102	136617	101192	237809	7	11	8	19	2	2	17
945	VMU130	262431	67296	329727	10	16	4	20	1	1	16
B46	VMU104	106466	13804	120270	4	6	1	7	0	0	14
B47	VMU106	26682	18437	45119	1	2	0	2	0	0	19
948	V1U113	23987	13370	37357	1	2	0	2	0	0	24
949	V1U110	51776	47327	99103	3	4	4	8	0	1	21
B4A	V2AIM2		158376	158376	5		7	7		0	14
B4B	V1U112	73903	12427	86330	3	4		4	0	0	16
94C											
94D											
B4E											
B4F											
x4x		1044023	541943	1585966	49	66	30	96			20
94x		715359	272180	987539	30	47	18	65			21
B4x		328664	269763	598427	19	19	12	31			17

KES 8/12/83 VMIO80T1

COMBINED I/O ACTIVITY, VM1 & VM2
Thursday June 30, 1983

CU ADD	CHAN ADD	VM1 I/O	VM2 I/O	TOTAL I/O	I/O SEC	CU BSY @3.5ms	CH BSY 4CHAN
82x		544980	487574	1032554	32	.112	
84x		988831	51999	1040830	32	.113	
86x		55845	1147328	1203173	37	.130	
	8xx	1589656	1686901	3276557	101		.089
92x		979251	251664	1230915	38	.133	
94x		715359	272180	987539	31	.107	
9Ax		439007	353800	792807	25	.086	
9Bx		294225	300558	594783	18	.064	
	9xx	2427842	1178202	3606044	111		.098
A2x		343989	213788	557777	17	.060	
A4x		835572	125504	961076	30	.104	
AAx		29438	198523	227961	7	.025	
ABx		203632	239351	442983	14	.048	
	Axx	1412631	777166	2189797	68		.059
B2x		431148	982305	1413453	44	.153	
B4x		328664	269763	598427	19	.065	
B6x		22291	497006	519297	16	.056	
	Bxx	782103	1749074	2531177	78		.068
	TOTALS	6212232	5391343	11603575	358		

76 Volumes Online 152678 I/Os per Volume

KES 8/12/83 VMIO80T1

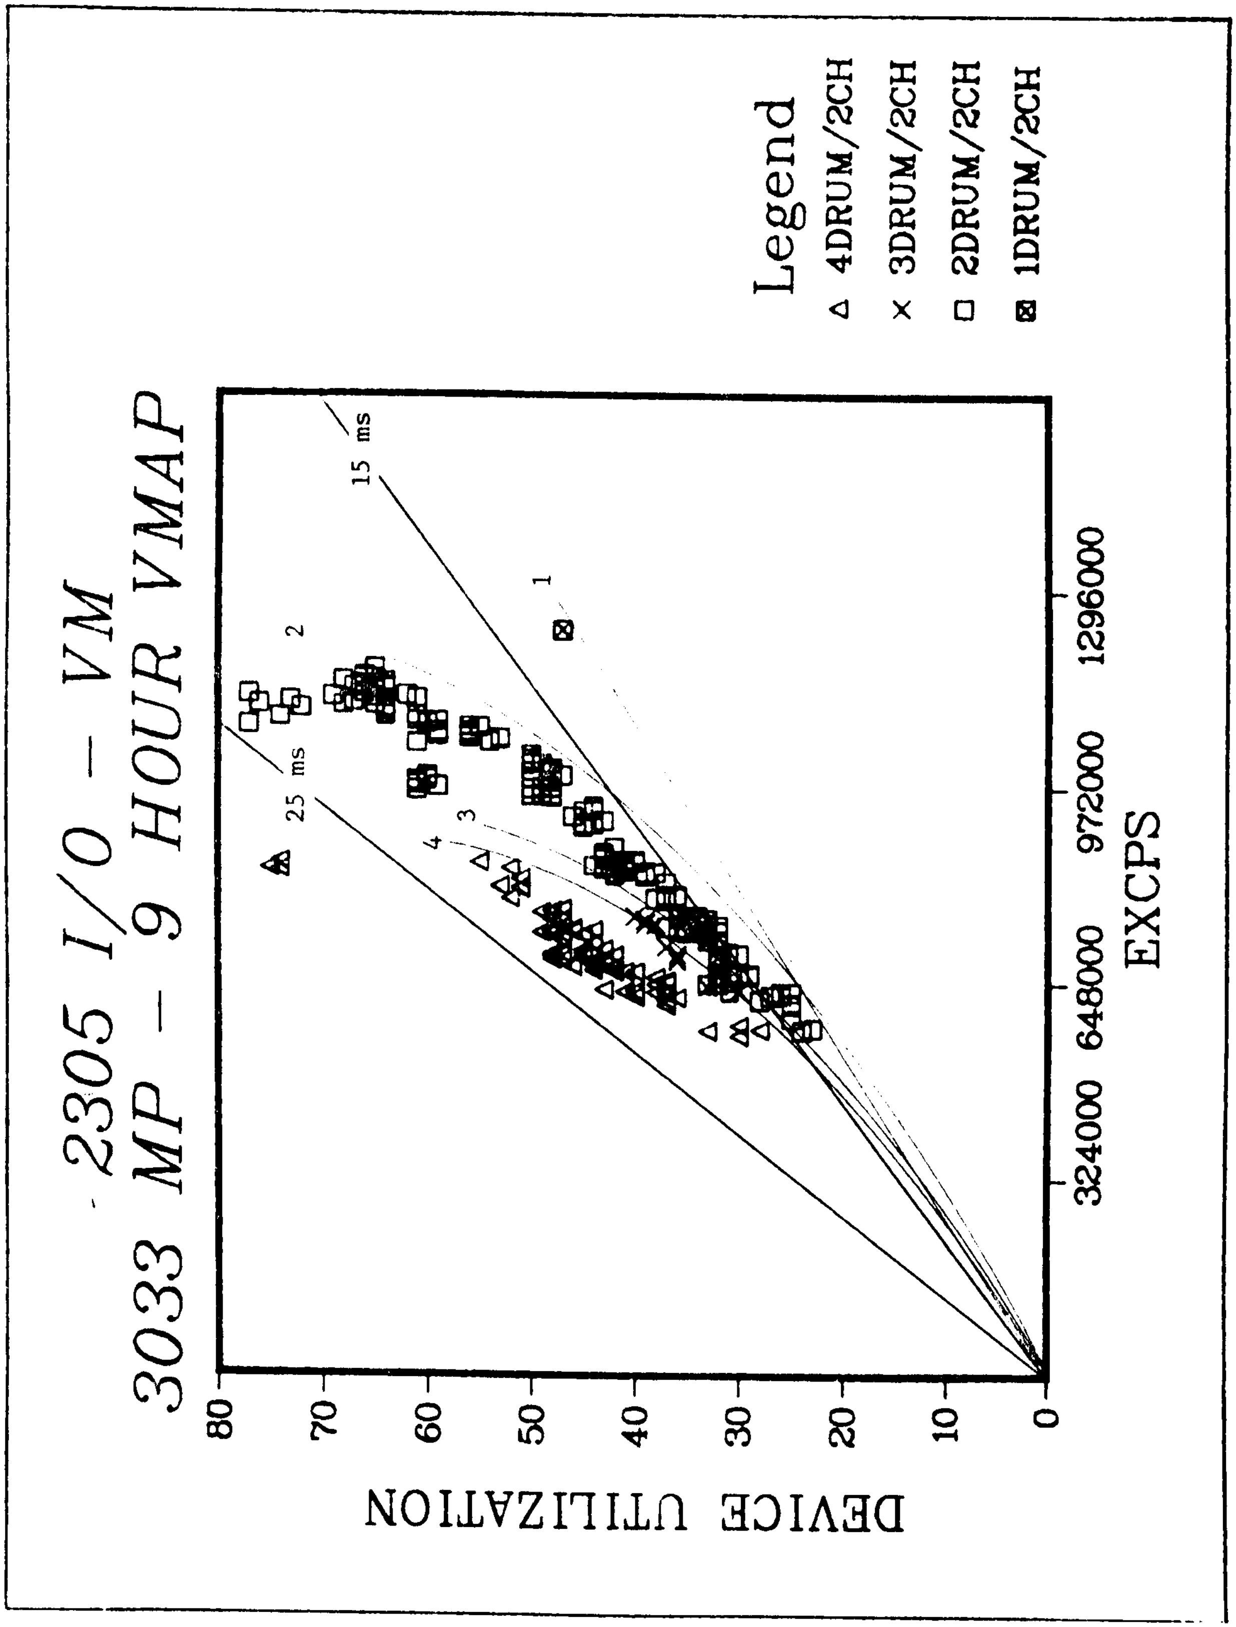
2305 I/O – VM
3033 MP – 9 HOUR VMAP
Legend
4DRUM/2CH
3DRUM/2CH
2DRUM/2CH
1DRUM/2CH
DEVICE UTILIZATION
0
10
20
30
40
50
60
70
80
EXCPS
324000
648000
972000
1296000
15 ms
25 ms
1
2
3
4

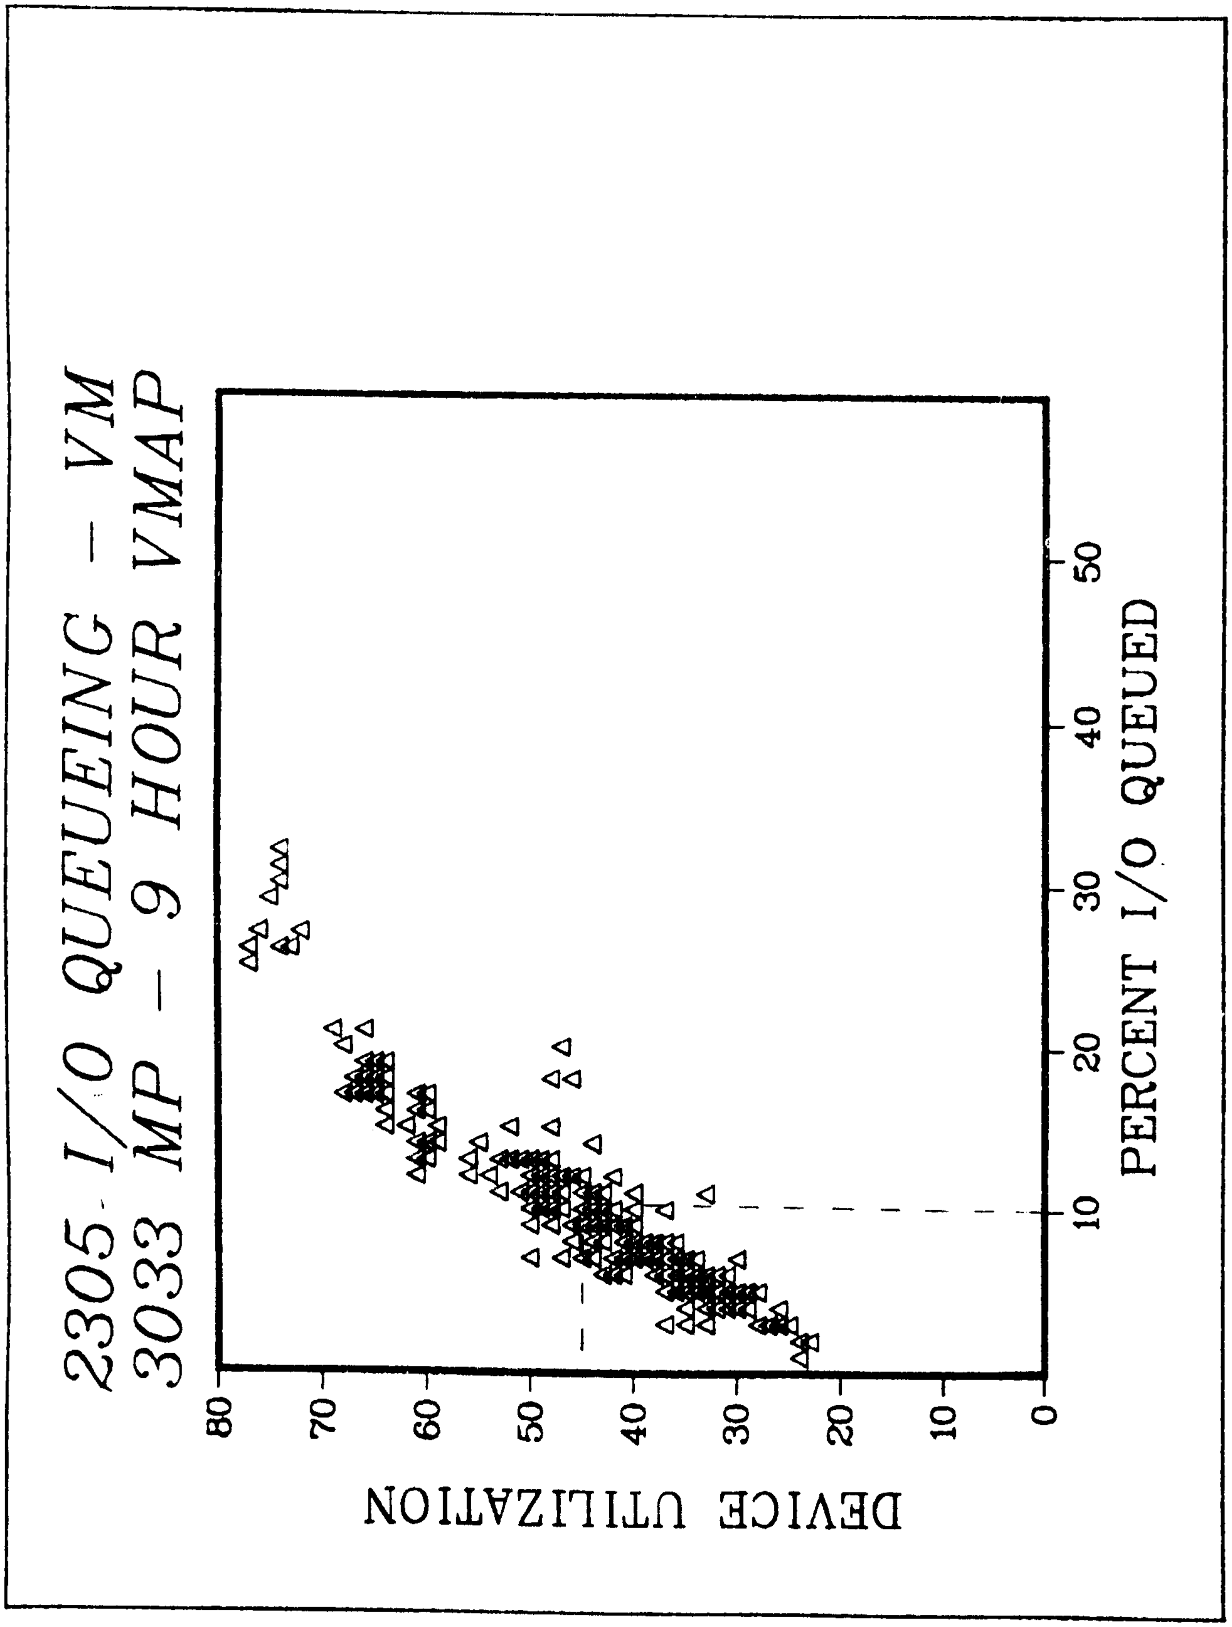
2305 I/O QUEUEING — VM
3033 MP — 9 HOUR VMAP
DEVICE UTILIZATION
PERCENT I/O QUEUED

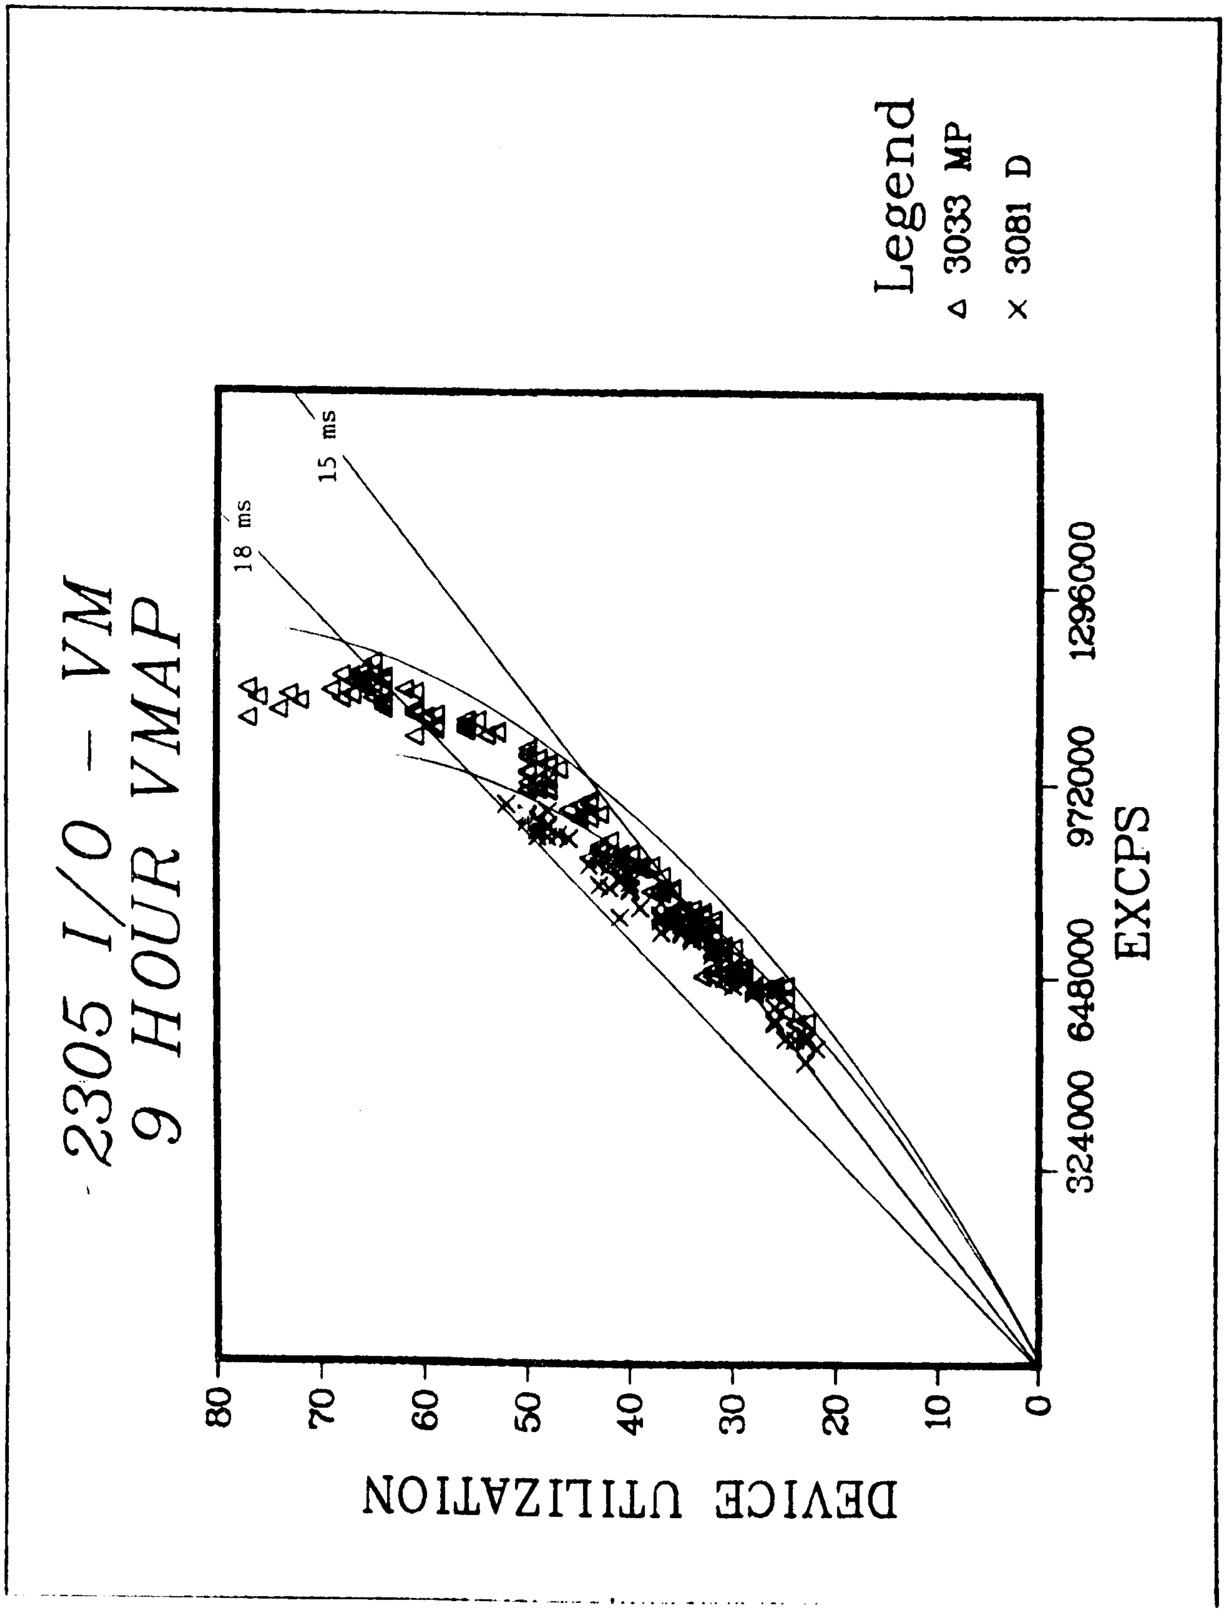
2305 I/O – VM
9 HOUR VMAP
Legend
△ 3033 MP
× 3081 D
18 ms
15 ms
DEVICE UTILIZATION
0
10
20
30
40
50
60
70
80
EXCPS
324000
648000
972000
1296000

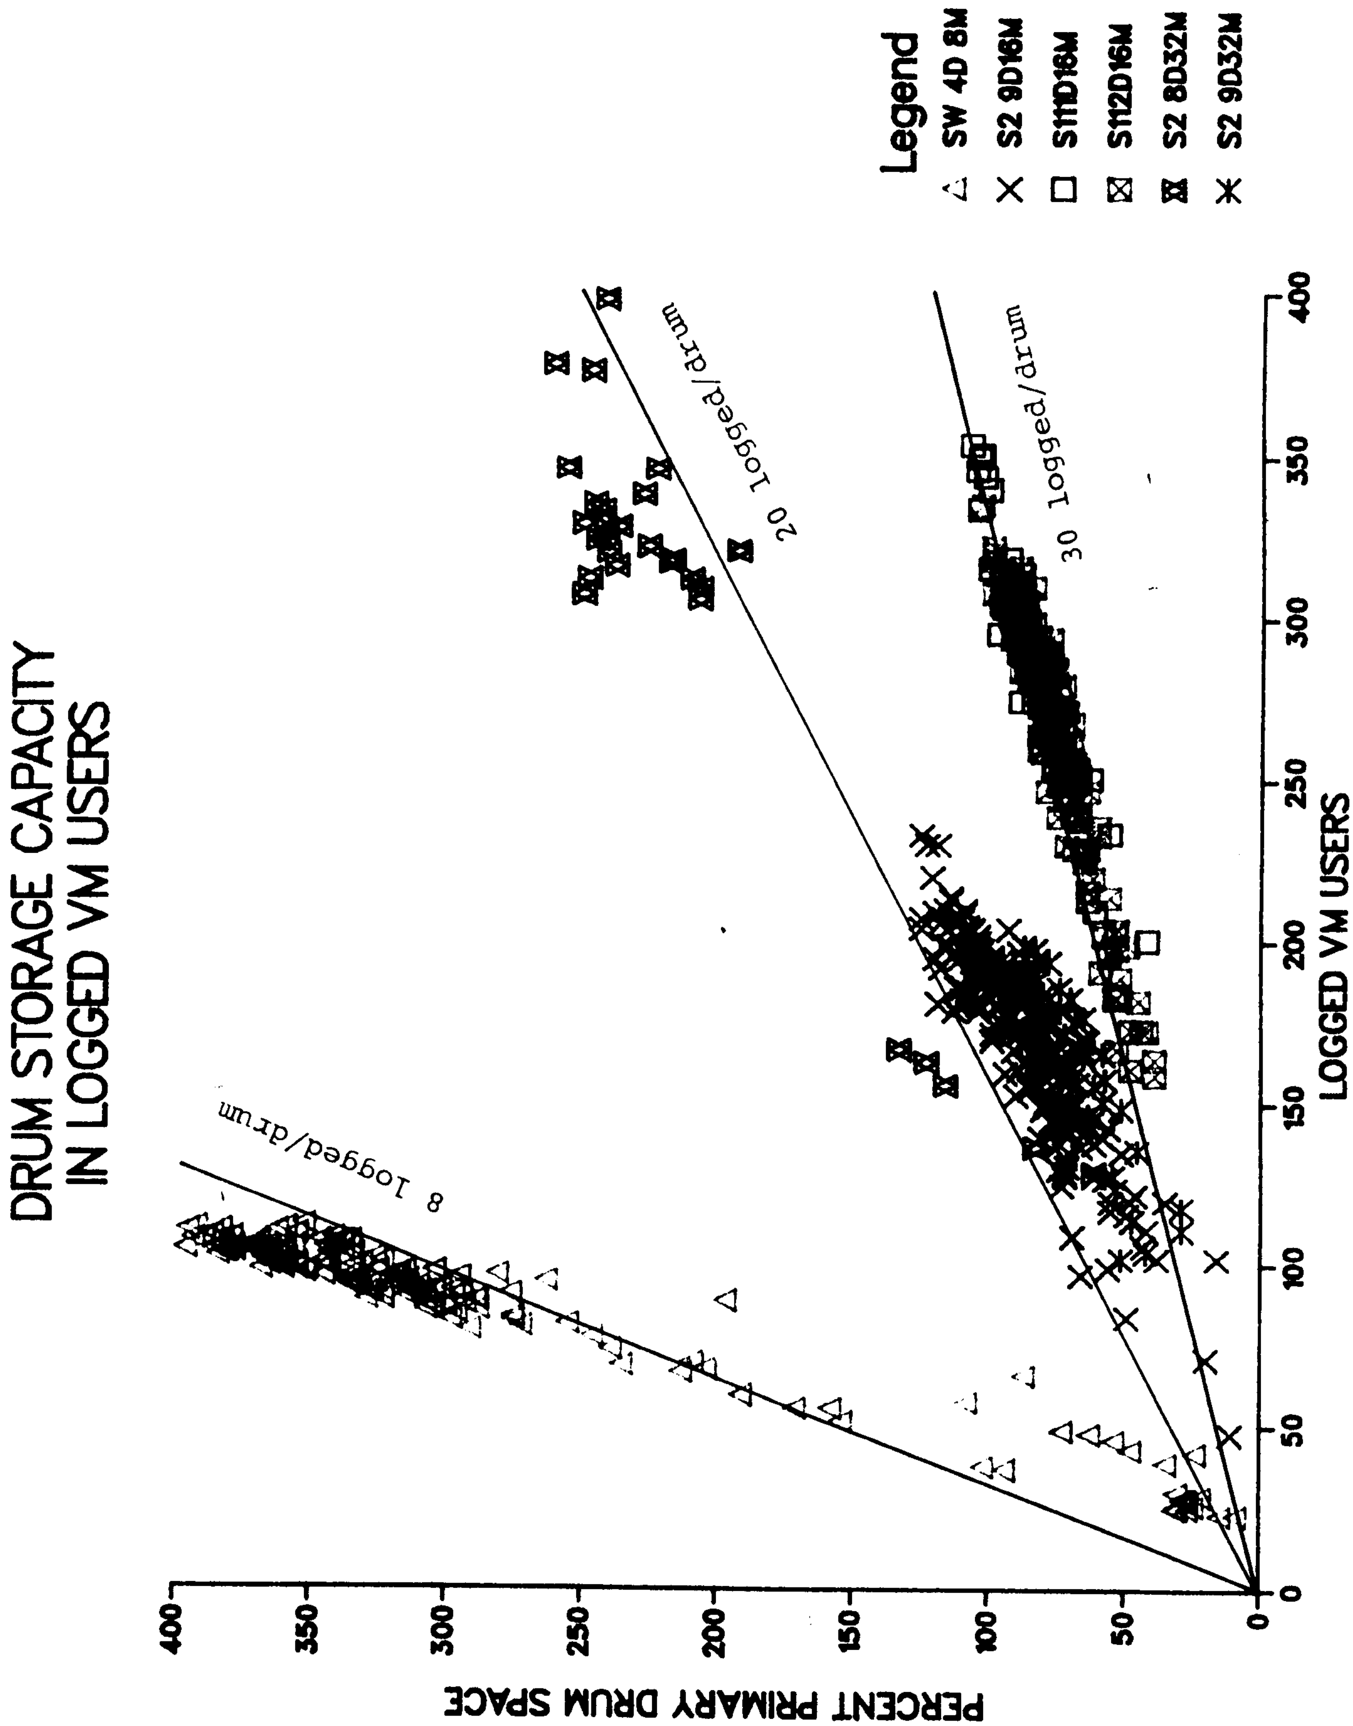
DRUM STORAGE CAPACITY
IN LOGGED VM USERS
Legend
SW 4D 8M
S2 9D16M
S11D16M
S112D16M
S2 8D32M
S2 9D32M
8 logged/drum
20 logged/drum
30 logged/drum
PERCENT PRIMARY DRUM SPACE
LOGGED VM USERS
0
50
100
150
200
250
300
350
400

CHANNEL UTILIZATION AND
3370 DASD PERFORMANCE

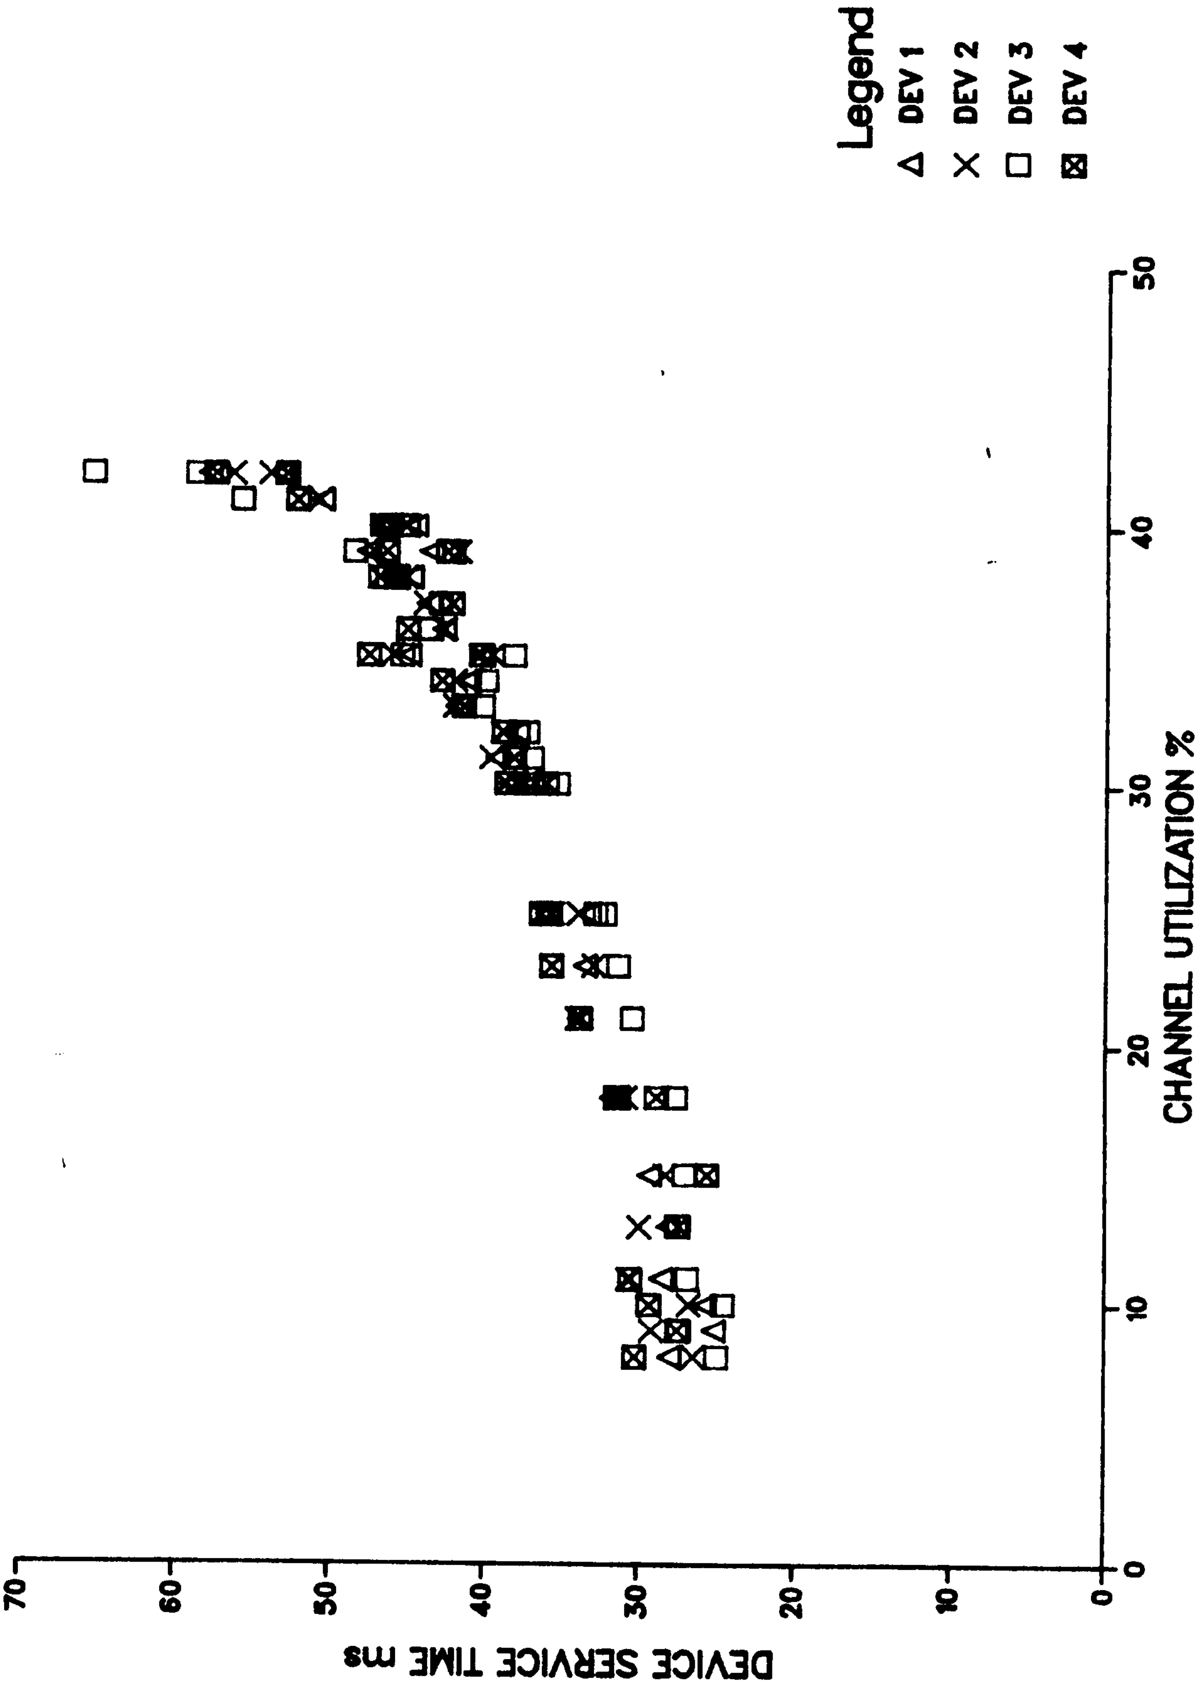

CHANNEL ACTIVITY RATE AND 3370 DASD PERFORMANCE

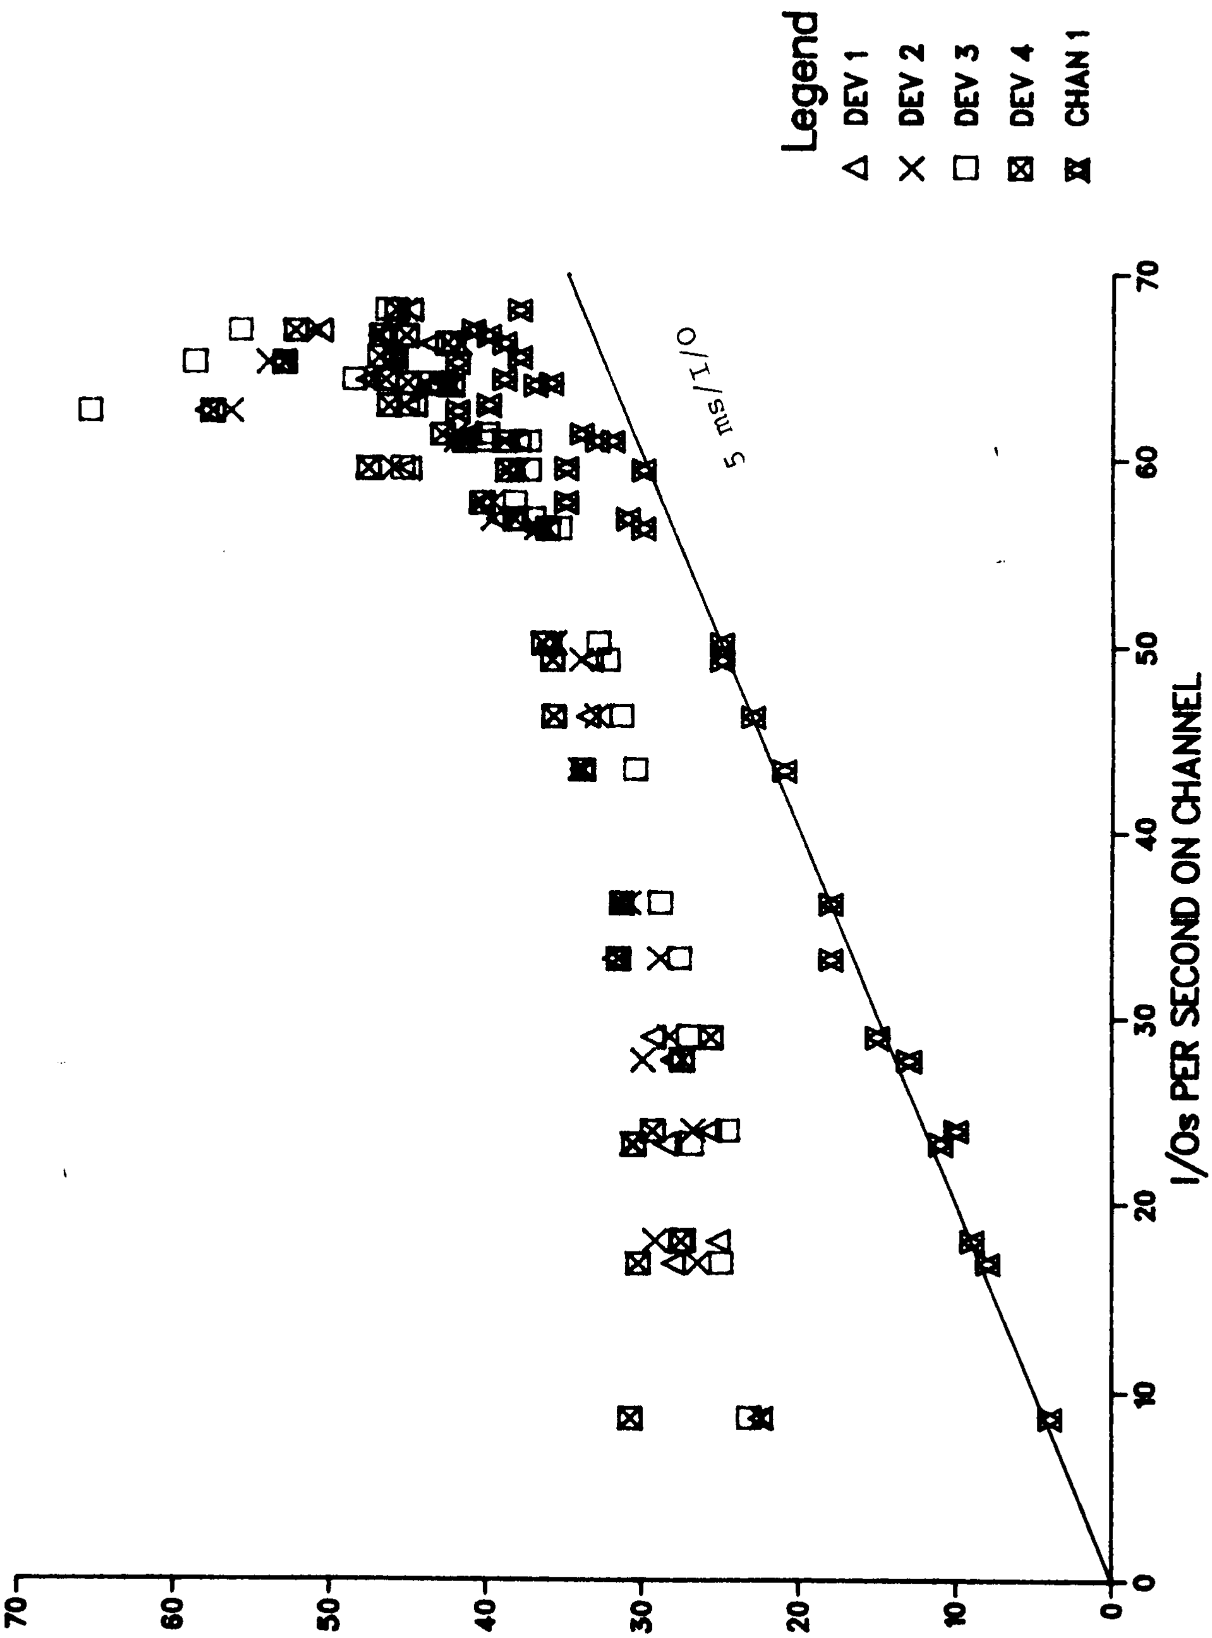

HIGH PERFORMANCE CMS INTENSIVE WORKLOAD
PERFORMANCE CHARACTERISTICS AND CONFIGURATION REQUIREMENTS

ARIP	ACTIVE	LOGGED	MEGS	PAGES	I/Os	DRUMS	DASD	CHAN
1	20	50	3	40	80	1	4	2+T&T
2	40	100	5	80	160	2	7	3+T&T
3	60	150	8	120	240	3	10	5+T&T
4	80	200	10	160	320	4	14	6+T&T
5	100	250	12	200	400	5	17	8+T&T
6	120	300	15	240	480	6	20	9+T&T
7	140	350	17	280	560	7	24	10+T&T
8	160	400	20	320	640	8	27	12+T&T
9	180	450	22	360	720	9	30	13+T&T
10	200	500	24	400	800	10	34	15+T&T
11	220	550	27	440	880	11	37	16+T&T
12	240	600	29	480	960	12	40	17+T&T
13	260	650	32	520	1040	13	44	19+T&T
14	280	700	34	560	1120	14	47	20+T&T
15	300	750	36	600	1200	15	50	22+T&T

ACTIVE = ARIP/.05 [50000 inst/act user/sec]

LOGGED = ACTIVE * 2.5

MEGS = (0.4 * ARIP) + (ACTIVE * WKSET [100])

PAGES = ACTIVE * WKSET/4 * 1.6 / THINKTIME [20 sec]

I/O = ARIP/.0125 [12500 inst / I/O]

DRUMS = PAGES/40

DASD = I/O * 0.5 / 12

CHAN = DRUMS/1 + DASD/8 + TAPE/? + TERMINAL/?

METHODOLOGY

IDENTIFY PERFORMANCE CONSTRAINTS

BALANCE SYSTEM RESOURCES

PLAN FUTURE RESOURCES

Sanitized Copy Approved for Release 2010/06/07 : CIA-RDP85-00142R000300200003-6

FUTURE

CONFIGURE & ANALYSE 3081 K – 32 MEGS

CHARACTERIZE SYSTEM RESOURCE REQUIREMENTS PER TERMINAL

PERFORMANCE ANALYST – EDUCATION, TRAINING, CAREER PATH, ...